



Citrix XenServer® 5.6 Administrator's Guide

Published June 2010
1.1 Edition



Citrix XenServer ® 5.6 Administrator's Guide

Copyright © 2009 Citrix All Rights Reserved.
Version: 5.6

Citrix, Inc.
851 West Cypress Creek Road
Fort Lauderdale, FL 33309
United States of America

Disclaimers

This document is furnished "AS IS." Citrix, Inc. disclaims all warranties regarding the contents of this document, including, but not limited to, implied warranties of merchantability and fitness for any particular purpose. This document may contain technical or other inaccuracies or typographical errors. Citrix, Inc. reserves the right to revise the information in this document at any time without notice. This document and the software described in this document constitute confidential information of Citrix, Inc. and its licensors, and are furnished under a license from Citrix, Inc.

Citrix Systems, Inc., the Citrix logo, Citrix XenServer and Citrix XenCenter, are trademarks of Citrix Systems, Inc. in the United States and other countries. All other products or services mentioned in this document are trademarks or registered trademarks of their respective companies.

Trademarks

Citrix ®
XenServer ®
XenCenter ®

1.1 Edition



Table of Contents

Document Overview	1
How this Guide relates to other documentation	1
Managing users	2
Authenticating users using Active Directory (AD)	2
Configuring Active Directory authentication	3
User authentication	5
Removing access for a user	6
Leaving an AD domain	6
Role Based Access Control	7
Roles	8
Definitions of RBAC roles and permissions	9
Working with RBAC using the xe CLI	14
To list all the available defined roles in XenServer	14
To display a list of current subjects:	15
To add a subject to RBAC	16
To assign an RBAC role to a created subject	16
To change a subject's RBAC role:	17
Auditing	17
Audit log xe CLI commands	17
To obtain all audit records from the pool	17
To obtain audit records of the pool since a precise millisecond timestamp	17
To obtain audit records of the pool since a precise minute timestamp	17
How does XenServer compute the roles for the session?	18
XenServer hosts and resource pools	19
Hosts and resource pools overview	19
Requirements for creating resource pools	19



Creating a resource pool	20
Creating heterogeneous resource pools	21
Adding shared storage	21
Removing a XenServer host from a resource pool	22
High Availability	23
HA Overview	23
Overcommitting	23
Overcommitment Warning	23
Host Fencing	23
Configuration Requirements	24
Restart priorities	24
Enabling HA on a XenServer pool	25
Enabling HA using the CLI	25
Removing HA protection from a VM using the CLI	26
Recovering an unreachable host	26
Shutting down a host when HA is enabled	26
Shutting down a VM when it is protected by HA	26
Host Power On	27
Powering on hosts remotely	27
Using the CLI to Manage Host Power On	27
To enable Host Power On using the CLI	28
To turn on hosts remotely using the CLI	28
Configuring a Custom Script for XenServer's Host Power On Feature	28
Key/Value Pairs	28
host.power_on_mode	28
host.power_on_config	29
Sample Script	29
Storage	30

Storage Overview	30
Storage Repositories (SRs)	30
Virtual Disk Images (VDIs)	30
Physical Block Devices (PBDs)	30
Virtual Block Devices (VBDs)	31
Summary of Storage objects	31
Virtual Disk Data Formats	31
VHD-based VDIs	32
VHD Chain Coalescing	32
Space Utilization	32
LUN-based VDIs	33
Storage configuration	33
Creating Storage Repositories	33
Upgrading LVM storage from XenServer 5.0 or earlier	34
LVM performance considerations	34
VDI types	34
Creating a raw virtual disk using the xe CLI	34
Converting between VDI formats	35
Probing an SR	35
Storage Multipathing	38
Storage Repository Types	39
Local LVM	40
Creating a local LVM SR (lvm)	40
Local EXT3 VHD	40
Creating a local EXT3 SR (ext)	40
udev	41
ISO	41
EqualLogic	41

Creating a shared EqualLogic SR	41
EqualLogic VDI Snapshot space allocation with XenServer EqualLogic Adapter	42
Creating a VDI using the CLI	43
NetApp	43
Creating a shared NetApp SR over iSCSI	46
Managing VDIs in a NetApp SR	47
Taking VDI snapshots with a NetApp SR	47
Software iSCSI Support	47
XenServer Host iSCSI configuration	48
Managing Hardware Host Bus Adapters (HBAs)	48
Sample QLogic iSCSI HBA setup	48
Removing HBA-based SAS, FC or iSCSI device entries	49
LVM over iSCSI	49
Creating a shared LVM over iSCSI SR using the software iSCSI initiator (lvmoiscsi)	49
Creating a shared LVM over Fibre Channel / iSCSI HBA or SAS SR (lvmohba)	50
NFS VHD	52
Creating a shared NFS SR (nfs)	53
LVM over hardware HBA	53
Citrix StorageLink Gateway (CSLG) SRs	54
Creating a shared StorageLink SR	54
Managing Storage Repositories	58
Destroying or forgetting a SR	59
Introducing an SR	59
Resizing an SR	60
Converting local Fibre Channel SRs to shared SRs	60
Moving Virtual Disk Images (VDIs) between SRs	60
Copying all of a VM's VDIs to a different SR	60
Copying individual VDIs to a different SR	60

Adjusting the disk IO scheduler	61
Virtual disk QoS settings	61
Configuring VM memory	63
What is Dynamic Memory Control (DMC)?	63
The concept of dynamic range	63
The concept of static range	64
DMC Behaviour	64
How does DMC Work?	64
Memory constraints	65
Supported operating systems	65
xen CLI commands	66
Display the static memory properties of a VM	66
Display the dynamic memory properties of a VM	66
Updating memory properties	67
Update individual memory properties	68
Upgrade issues	68
Workload Balancing interaction	68
Networking	69
XenServer networking overview	69
Network objects	70
Networks	70
VLANs	70
Using VLANs with host management interfaces	70
Using VLANs with virtual machines	70
Using VLANs with dedicated storage NICs	71
Combining management interfaces and guest VLANs on a single host NIC	71
NIC bonds	71

Initial networking configuration	72
Managing networking configuration	72
Creating networks in a standalone server	72
Creating networks in resource pools	73
Creating VLANs	73
Creating NIC bonds on a standalone host	74
Creating a NIC bond on a dual-NIC host	74
Controlling the MAC address of the bond	75
Reverting NIC bonds	75
Creating NIC bonds in resource pools	75
Adding NIC bonds to new resource pools	76
Adding NIC bonds to an existing pool	77
Configuring a dedicated storage NIC	79
Controlling Quality of Service (QoS)	79
Changing networking configuration options	79
Hostname	80
DNS servers	80
Changing IP address configuration for a standalone host	80
Changing IP address configuration in resource pools	80
Management interface	81
Disabling management access	81
Adding a new physical NIC	82
NIC/PIF ordering in resource pools	82
Verifying NIC ordering	82
Re-ordering NICs	82
Networking Troubleshooting	83
Diagnosing network corruption	83
Recovering from a bad network configuration	84



Workload Balancing 85

What's New?	85
New Features	85
Changes	86
Workload Balancing Overview	86
Workload Balancing Basic Concepts	87
Workload Balancing Installation Overview	87
Workload Balancing System Requirements	88
Supported XenServer Versions	88
Supported Operating Systems	88
Recommended Hardware	88
Workload Balancing Data Store Requirements	89
SQL Server Database Authentication Requirements	89
Operating System Language Support	90
Preinstallation Considerations	90
WLB Access Control Permissions	91
Installing Workload Balancing	91
To install Workload Balancing server	92
To verify your Workload Balancing installation	93
Configuring Firewalls	94
Upgrading Workload Balancing	94
Upgrading Workload Balancing on the Same Operating System	95
Upgrading SQL Server	95
Upgrading Workload Balancing and the Operating System	95
Initializing Workload Balancing	95
To initialize Workload Balancing	96
Authorization for Workload Balancing	97
Configuring Antivirus Software	98

Configuring Workload Balancing Settings	98
To display the Workload Balancing Configuration dialog box	99
Adjusting the Optimization Mode	99
Fixed	99
Scheduled	99
To set an optimization mode for all time periods	100
To specify times when the optimization mode will change automatically	100
To edit or delete an automatic optimization interval	100
Optimizing and Managing Power Automatically	100
Accepting Optimization Recommendations Automatically	101
Enabling Workload Balancing Power Management	101
Designing Environments for Power Management and VM Consolidation	102
To apply optimization recommendations automatically	103
To select servers for power management	103
Changing the Critical Thresholds	103
Default Settings for Critical Thresholds	104
To change the critical thresholds	104
Tuning Metric Weightings	104
To edit metric weighting factors	105
Excluding Hosts from Recommendations	105
To exclude hosts from placement and optimization recommendations	105
Configuring Optimization Intervals, Report Subscriptions, and Data Storage	106
Historical Data (Storage Time)	106
To configure the data storage period	106
VM Optimization Criteria	106
Length of Time Between Optimization Recommendations After VM Moves	106
Number of Times an Optimization Recommendation is Made	107
Setting the Minimum Optimization Severity	107

Modifying the Aggressiveness Setting	107
Receiving Reports by Email Automatically (Report Subscriptions)	108
To configure report subscriptions	109
Choosing an Optimal Server for VM Initial Placement, Migrate, and Resume	109
To start a virtual machine on the optimal server	109
To resume a virtual machine on the optimal server	109
Accepting Optimization Recommendations	110
To accept an optimization recommendation	110
Administering Workload Balancing	110
Disabling Workload Balancing	111
Reconfiguring a Pool to Use Another WLB Server	111
Updating Workload Balancing Credentials	112
Uninstalling Workload Balancing	113
Customizing Workload Balancing	113
Entering Maintenance Mode with Workload Balancing Enabled	113
To enter maintenance mode with Workload Balancing enabled	114
Working with Workload Balancing Reports	114
Introduction	114
Subscribing to Workload Balancing Reports	114
Using Workload Balancing Reports for Tasks	115
Evaluating the Effectiveness of Your Optimization Thresholds	115
Generating and Managing Workload Balancing Reports	115
To generate a Workload Balancing report	115
To subscribe to a Workload Balancing report	115
To cancel a report subscription	116
To navigate in a Workload Balancing Report	116
To print a Workload Balancing report	117
To export a Workload Balancing report	117

Displaying Workload Balancing Reports	117
Report Generation Features	117
Toolbar Buttons	117
Workload Balancing Report Glossary	118
Host Health History	118
Pool Optimization Performance History	119
Pool Audit Log History	119
Audit Log Event Names	120
Pool Health	120
Pool Health History	121
Pool Optimization History	121
Virtual Machine Motion History	122
Virtual Machine Performance History	122
Backup and recovery	123
Backups	123
Full metadata backup and disaster recovery (DR)	124
DR and metadata backup overview	124
Backup and restore using xsconsole	124
Moving SRs between hosts and Pools	125
Using Portable SRs for Manual Multi-Site Disaster Recovery	126
VM Snapshots	126
Regular Snapshots	127
Quiesced Snapshots	127
Snapshots with memory	127
Creating a VM Snapshot	127
Creating a snapshot with memory	128
To list all of the snapshots on a XenServer pool	128
To list the snapshots on a particular VM	128

Restoring a VM to its previous state	129
Deleting a snapshot	129
Snapshot Templates	130
Creating a template from a snapshot	130
Exporting a snapshot to a template	131
Advanced Notes for Quiesced Snapshots	131
Coping with machine failures	133
Member failures	133
Master failures	133
Pool failures	134
Coping with Failure due to Configuration Errors	134
Physical Machine failure	134
Monitoring and managing XenServer	136
Alerts	136
Customizing Alerts	137
Configuring Email Alerts	138
Custom Fields and Tags	139
Custom Searches	139
Determining throughput of physical bus adapters	139
Troubleshooting	140
XenServer host logs	140
Sending host log messages to a central server	140
XenCenter logs	141
Troubleshooting connections between XenCenter and the XenServer host	141
A. Command line interface	142
Basic xe syntax	142
Special characters and syntax	143

Command types	144
Parameter types	145
Low-level param commands	145
Low-level list commands	146
xe command reference	147
Bonding commands	147
bond-create	147
bond-destroy	147
CD commands	148
cd-list	149
Console commands	149
Event commands	150
event-wait	150
Host (XenServer host) commands	151
host-backup	154
host-bugreport-upload	155
host-crashdump-destroy	155
host-crashdump-upload	155
host-disable	155
host-dmesg	155
host-emergency-management-reconfigure	155
host-enable	156
host-evacuate	156
host-forget	156
host-get-system-status	156
host-get-system-status-capabilities	157
host-is-in-emergency-mode	158
host-apply-edition	158

license-server-address	158
license-server-port	158
host-license-add	158
host-license-view	159
host-logs-download	159
host-management-disable	159
host-management-reconfigure	159
host-power-on	160
host-set-power-on	160
host-reboot	160
host-restore	160
host-set-hostname-live	160
host-shutdown	161
host-syslog-reconfigure	161
Log commands	161
log-get-keys	161
log-reopen	161
log-set-output	161
Message commands	162
message-create	162
message-list	162
Network commands	163
network-create	164
network-destroy	164
Patch (update) commands	164
patch-apply	165
patch-clean	165
patch-pool-apply	165

patch-precheck	165
patch-upload	165
PBD commands	165
pbd-create	166
pbd-destroy	166
pbd-plug	166
pbd-unplug	166
PIF commands	166
pif-forget	169
pif-introduce	169
pif-plug	170
pif-reconfigure-ip	170
pif-scan	170
pif-unplug	170
Pool commands	170
pool-designate-new-master	172
pool-dump-database	172
pool-eject	172
pool-emergency-reset-master	172
pool-emergency-transition-to-master	172
pool-ha-enable	172
pool-ha-disable	173
pool-join	173
pool-recover-slaves	173
pool-restore-database	173
pool-sync-database	173
Storage Manager commands	173
SR commands	174

sr-create	175
sr-destroy	175
sr-forget	175
sr-introduce	176
sr-probe	176
sr-scan	176
Task commands	176
task-cancel	177
Template commands	177
template-export	185
Update commands	185
update-upload	185
User commands	185
user-password-change	185
VBD commands	186
vbd-create	187
vbd-destroy	188
vbd-eject	188
vbd-insert	188
vbd-plug	188
vbd-unplug	188
VDI commands	188
vdi-clone	190
vdi-copy	190
vdi-create	190
vdi-destroy	191
vdi-forget	191
vdi-import	191

vdi-introduce	191
vdi-resize	191
vdi-snapshot	191
vdi-unlock	192
VIF commands	192
vif-create	194
vif-destroy	194
vif-plug	194
vif-unplug	194
VLAN commands	195
vlan-create	195
pool-vlan-create	195
vlan-destroy	195
VM commands	195
vm-cd-add	202
vm-cd-eject	202
vm-cd-insert	203
vm-cd-list	203
vm-cd-remove	203
vm-clone	203
vm-compute-maximum-memory	203
vm-copy	204
vm-crashdump-list	204
vm-data-source-forget	204
vm-data-source-list	204
vm-data-source-query	205
vm-data-source-record	205
vm-destroy	205

vm-disk-add	205
vm-disk-list	206
vm-disk-remove	206
vm-export	206
vm-import	206
vm-install	207
vm-memory-shadow-multiplier-set	207
vm-migrate	207
vm-reboot	208
vm-reset-powerstate	208
vm-resume	208
vm-shutdown	208
vm-start	209
vm-suspend	209
vm-uninstall	209
vm-vcpu-hotplug	209
vm-vif-list	210
Workload Balancing commands	210
pool-initialize-wlb	210
pool-param-set other-config	210
host-retrieve-wlb-evacuate-recommendations	210
vm-retrieve-wlb-recommendations	210
pool-certificate-list	211
pool-certificate-install	211
pool-certificate-sync	211
pool-param-set	212
pool-deconfigure-wlb	212
pool-retrieve-wlb-configuration	212



pool-retrieve-wlb-recommendations	212
pool-retrieve-wlb-report	212
pool-send-wlb-configuration	213
Index	214



Document Overview

This document is a system administrator's guide to XenServer™, the platform virtualization solution from Citrix®. It describes the tasks involved in configuring a XenServer deployment-- in particular, how to set up storage, networking and resource pools, and how to administer XenServer hosts using the xe command line interface (CLI).

This section summarizes the rest of the guide so that you can find the information you need. The following topics are covered:

- XenServer hosts and resource pools
- XenServer storage configuration
- XenServer network configuration
- XenServer workload balancing
- XenServer backup and recovery
- Monitoring and managing XenServer
- XenServer command line interface
- XenServer troubleshooting
- XenServer resource allocation guidelines

How this Guide relates to other documentation

This document is primarily aimed at system administrators, who need to configure and administer XenServer deployments. Other documentation shipped with this release includes:

- *XenServer Installation Guide* provides a high level overview of XenServer, along with step-by-step instructions on installing XenServer hosts and the XenCenter management console.
- *XenServer Virtual Machine Installation Guide* describes how to install Linux and Windows VMs on top of a XenServer deployment. As well as installing new VMs from install media (or using the VM *templates* provided with the XenServer release), this guide also explains how to create VMs from existing physical machines, using a process called *P2V*.
- *XenServer Software Development Kit Guide* presents an overview of the XenServer SDK- a selection of code samples that demonstrate how to write applications that interface with XenServer hosts.
- *XenAPI Specification* provides a programmer's reference guide to the XenServer API.
- *XenServer User Security* considers the issues involved in keeping your XenServer installation secure.
- *Release Notes* provides a list of known issues that affect this release.

Managing users

When you first install XenServer, a user account is added to XenServer automatically. This account is the local super user (LSU), or root, which is authenticated locally by the XenServer computer.

The *local super user (LSU)*, or root, is a special user account used for system administration and has all rights or permissions. In XenServer, the local super user is the default account at installation. The LSU is authenticated by XenServer and not an external authentication service. This means that if the external authentication service fails, the LSU can still log in and manage the system. The LSU can always access the XenServer physical server through SSH.

You can create additional users by adding their Active Directory accounts through either the XenCenter's Users tab or the CLI. All editions of XenServer can add user accounts from Active Directory. However, only XenServer Enterprise and Platinum editions let you assign these Active Directory accounts different levels of permissions (through the Role Based Access Control (RBAC) feature). If you do not use Active Directory in your environment, you are limited to the LSU account.

The permissions assigned to users when you first add their accounts varies according to your version of XenServer:

- In the XenServer and XenServer Advanced edition, when you create (add) new users, XenServer automatically grants the accounts access to all features available in that version.
- In the XenServer Enterprise and Platinum editions, when you create new users, XenServer does not assign newly created user accounts roles automatically. As a result, these accounts do not have any access to the XenServer pool until you assign them a role.

If you do not have one of these editions, you can add users from Active Directory. However, all users will have the Pool Administrator role.

These permissions are granted through roles, as discussed in [the section called “Authenticating users using Active Directory \(AD\)”](#).

Authenticating users using Active Directory (AD)

If you want to have multiple user accounts on a server or a pool, you must use Active Directory user accounts for authentication. This lets XenServer users log in to a pool's XenServers using their Windows domain credentials.

The only way you can configure varying levels of access for specific users is by enabling Active Directory authentication, adding user accounts, and assign roles to those accounts.

Active Directory users can use the xe CLI (passing appropriate `-u` and `-pw` arguments) and also connect to the host using XenCenter. Authentication is done on a per-resource pool basis.

Access is controlled by the use of *subjects*. A subject in XenServer maps to an entity on your directory server (either a user or a group). When external authentication is enabled, the credentials used to create a session are first checked against the local root credentials (in case your directory server is unavailable) and then against the subject list. To permit access, you must create a subject entry for the person or group you wish to grant access to. This can be done using XenCenter or the xe CLI.

If you are familiar with XenCenter, note that the XenServer CLI uses slightly different terminology to refer to Active Directory and user account features:

XenCenter Term	XenServer CLI Term
Users	Subjects
Add users	Add subjects

Understanding Active Directory authentication in the XenServer environment

Even though XenServers are Linux-based, XenServer lets you use Active Directory accounts for XenServer user accounts. To do so, it passes Active Directory credentials to the Active Directory domain controller.

When added to XenServer, Active Directory users and groups become XenServer subjects, generally referred to as simply users in XenCenter. When a subject is registered with XenServer, users/groups are authenticated with Active Directory on login and do not need to qualify their user name with a domain name.

Note:

By default, if you did not qualify the user name (for example, enter either `mydomain\myuser` or `myuser@mydomain.com`), XenCenter always attempts to log users in to Active Directory authentication servers using the domain to which it is currently joined. The exception to this is the LSU account, which XenCenter always authenticates locally (that is, on the XenServer) first.

The external authentication process works as follows:

1. The credentials supplied when connecting to a server are passed to the Active Directory domain controller for authentication.
2. The domain controller checks the credentials. If they are invalid, the authentication fails immediately.
3. If the credentials are valid, the Active Directory controller is queried to get the subject identifier and group membership associated with the credentials.
4. If the subject identifier matches the one stored in the XenServer, the authentication is completed successfully.

When you join a domain, you enable Active Directory authentication for the pool. However, when a pool is joined to a domain, only users in that domain (or a domain with which it has trust relationships) can connect to the pool.

Note:

Manually updating the DNS configuration of a DHCP-configured network PIF is unsupported and might cause Active Directory integration, and consequently user authentication, to fail or stop working.

Upgrading from XenServer 5.5

When you upgrade from XenServer 5.5 to the current release, any user accounts created in XenServer 5.5 are assigned the role of pool-admin. This is done for backwards compatibility reasons: in XenServer 5.5, all users had full permissions to perform any task on the pool.

As a result, if you are upgrading from XenServer 5.5, make sure you revisit the role associated with each user account to make sure it is still appropriate.

Configuring Active Directory authentication

XenServer supports use of Active Directory servers using Windows 2003 or later.

Active Directory authentication for a XenServer host requires that the same DNS servers are used for both the Active Directory server (configured to allow for interoperability) and the XenServer host. In some



configurations, the active directory server may provide the DNS itself. This can be achieved either using DHCP to provide the IP address and a list of DNS servers to the XenServer host, or by setting values in the PIF objects or using the installer if a manual static configuration is used.

Citrix recommends enabling DHCP to broadcast host names. In particular, the host names `localhost` or `linux` should not be assigned to hosts.

Note the following:

- XenServer hostnames should be unique throughout the XenServer deployment. XenServer labels its AD entry on the AD database using its hostname. Therefore, if two XenServer hosts have the same hostname and are joined to the same AD domain, the second XenServer will overwrite the AD entry of the first XenServer, regardless of if they are in the same or in different pools, causing the AD authentication on the first XenServer to stop working.

It is possible to use the same hostname in two XenServer hosts, as long as they join different AD domains.

- The servers can be in different time-zones, as it is the UTC time that is compared. To ensure synchronization is correct, you may choose to use the same NTP servers for your XenServer pool and the Active Directory server.
- Mixed-authentication pools are not supported (that is, you cannot have a pool where some servers in the pool are configured to use Active Directory and some are not).
- The XenServer Active Directory integration uses the Kerberos protocol to communicate with the Active Directory servers. Consequently, XenServer does not support communicating with Active Directory servers that do not utilize Kerberos.
- For external authentication using Active Directory to be successful, it is important that the clocks on your XenServer hosts are synchronized with those on your Active Directory server. When XenServer joins the Active Directory domain, this will be checked and authentication will fail if there is too much skew between the servers.

Warning:

Host names must consist solely of no more than 63 alphanumeric characters, and must not be purely numeric.

Once you have Active Directory authentication enabled, if you subsequently add a server to that pool, you are prompted to configure Active Directory on the server joining the pool. When you are prompted for credentials on the joining server, enter Active Directory credentials with sufficient privileges to add servers to that domain.

Enabling external authentication on a pool

- External authentication using Active Directory can be configured using either XenCenter or the CLI using the command below.

```
xe pool-enable-external-auth auth-type=AD \  
  service-name=<full-qualified-domain> \  
  config:user=<username> \  
  config:pass=<password>
```

The user specified needs to have `Add/remove computer objects or workstations` privileges, which is the default for domain administrators.

Note:

If you are not using DHCP on the network that Active Directory and your XenServer hosts use you can use these two approaches to setup your DNS:

1. Configure the DNS server to use on your XenServer hosts:

```
xe pif-reconfigure-ip mode=static dns=<dnshost>
```

2. Manually set the management interface to use a PIF that is on the same network as your DNS server:

```
xe host-management-reconfigure pif-uuid=<pif_in_the_dns_subnetwork>
```

Note:

External authentication is a per-host property. However, Citrix advises that you enable and disable this on a per-pool basis – in this case XenServer will deal with any failures that occur when enabling authentication on a particular host and perform any roll-back of changes that may be required, ensuring that a consistent configuration is used across the pool. Use the **host-param-list** command to inspect properties of a host and to determine the status of external authentication by checking the values of the relevant fields.

Disabling external authentication

- Use XenCenter to disable Active Directory authentication, or the following xe command:

```
xe pool-disable-external-auth
```

User authentication

To allow a user access to your XenServer host, you must add a subject for that user or a group that they are in. (Transitive group memberships are also checked in the normal way, for example: adding a subject for group A, where group A contains group B and `user 1` is a member of group B would permit access to `user 1`.) If you wish to manage user permissions in Active Directory, you could create a single group that you then add and remove users to/from; alternatively, you can add and remove individual users from XenServer, or a combination of users and groups as your would be appropriate for your authentication requirements. The subject list can be managed from XenCenter or using the CLI as described below.

When authenticating a user, the credentials are first checked against the local root account, allowing you to recover a system whose AD server has failed. If the credentials (i.e.. username then password) do not match/authenticate, then an authentication request is made to the AD server – if this is successful the user's information will be retrieved and validated against the local subject list, otherwise access will be denied. Validation against the subject list will succeed if the user or a group in the transitive group membership of the user is in the subject list.

Note:

When using Active Directory groups to grant access for Pool Administrator users who will require host ssh access, the number of users in the Active Directory group must not exceed 500.

Allowing a user access to XenServer using the CLI

- To add an AD subject to XenServer:

```
xe subject-add subject-name=<entity name>
```

The entity name should be the name of the user or group to which you want to grant access. You may optionally include the domain of the entity (for example, '`<xend\user1>`' as opposed to '`<user1>`') although the behavior will be the same unless disambiguation is required.

Removing access for a user using the CLI

1. Identify the subject identifier for the subject you wish to revoke access. This would be the user or the group containing the user (removing a group would remove access to all users in that group, providing they are not also specified in the subject list). You can do this using the subject list command:



```
xe subject-list
```

You may wish to apply a filter to the list, for example to get the subject identifier for a user named `user1` in the `testad` domain, you could use the following command:

```
xe subject-list other-config:subject-name='<domain\user>'
```

2. Remove the user using the **subject-remove** command, passing in the subject identifier you learned in the previous step:

```
xe subject-remove subject-identifier=<subject identifier>
```

3. You may wish to terminate any current session this user has already authenticated. See [Terminating all authenticated sessions using xe](#) and [Terminating individual user sessions using xe](#) for more information about terminating sessions. If you do not terminate sessions the users whose permissions have been revoked may be able to continue to access the system until they log out.

Listing subjects with access

- To identify the list of users and groups with permission to access your XenServer host or pool, use the following command:

```
xe subject-list
```

Removing access for a user

Once a user is authenticated, they will have access to the server until they end their session, or another user terminates their session. Removing a user from the subject list, or removing them from a group that is in the subject list, will not automatically revoke any already-authenticated sessions that the user has; this means that they may be able to continue to access the pool using XenCenter or other API sessions that they have already created. In order to terminate these sessions forcefully, XenCenter and the CLI provide facilities to terminate individual sessions, or all currently active sessions. See the XenCenter help for more information on procedures using XenCenter, or below for procedures using the CLI.

Terminating all authenticated sessions using xe

- Execute the following CLI command:

```
xe session-subject-identifier-logout-all
```

Terminating individual user sessions using xe

1. Determine the subject identifier whose session you wish to log out. Use either the **session-subject-identifier-list** or **subject-list** xe commands to find this (the first shows users who have sessions, the second shows all users but can be filtered, for example, using a command like **xe subject-list other-config:subject-name=xendt\\user1** – depending on your shell you may need a double-backslash as shown).
2. Use the **session-subject-logout** command, passing the subject identifier you have determined in the previous step as a parameter, for example:

```
xe session-subject-identifier-logout subject-identifier=<subject-id>
```

Leaving an AD domain

Warning:



When you leave the domain (that is, disable Active Directory authentication and disconnect a pool or server from its domain), any users who authenticated to the pool or server with Active Directory credentials are disconnected.

Use XenCenter to leave an AD domain. See the XenCenter help for more information. Alternately run the **pool-disable-external-auth** command, specifying the pool uuid if required.

Note:

Leaving the domain will not cause the host objects to be removed from the AD database. See [this](#) knowledge base article for more information about this and how to remove the disabled host entries.

Role Based Access Control

Note:


The full RBAC feature is only available in **Citrix XenServer Enterprise Edition** or higher. To learn more about upgrading XenServer, click [here](#).

XenServer's Role Based Access Control (RBAC) allows you to assign users, roles, and permissions to control who has access to your XenServer and what actions they can perform. The XenServer RBAC system maps a user (or a group of users) to defined roles (a named set of permissions), which in turn have associated XenServer permissions (the ability to perform certain operations).

As users are not assigned permissions directly, but acquire them through their assigned role, management of individual user permissions becomes a matter of simply assigning the user to the appropriate role; this simplifies common operations. XenServer maintains a list of authorized users and their roles.

RBAC allows you to easily restrict which operations different groups of users can perform - thus reducing the probability of an accident by an inexperienced user.

To facilitate compliance and auditing, RBAC also provides an Audit Log feature and its corresponding Workload Balancing Pool Audit Trail report.



RBAC depends on Active Directory for authentication services. Specifically, XenServer keeps a list of authorized users based on Active Directory user and group accounts. As a result, you must join the pool to the domain and add Active Directory accounts before you can assign roles.



The local super user (LSU), or root, is a special user account used for system administration and has all rights or permissions. In XenServer, the local super user is the default account at installation. The LSU is authenticated via XenServer and not external authentication service, so if the external authentication service fails, the LSU can still log in and manage the system. The LSU can always access the XenServer physical host via SSH.

RBAC process

This is the standard process for implementing RBAC and assigning a user or group a role:

1. Join the domain. See [Enabling external authentication on a pool](#)
2. Add an Active Directory user or group to the pool. This becomes a subject. See [the section called "To add a subject to RBAC"](#).
3. Assign (or modify) the subject's RBAC role. See [the section called "To assign an RBAC role to a created subject"](#).

Roles

XenServer is shipped with the following six, pre-established roles:

- *Pool Administrator* (Pool Admin) – the same as being the local root. Can perform all operations.

Note:

The local super user (root) will always have the "Pool Admin" role. The Pool Admin role has the same permissions as the local root.

- *Pool Operator* (Pool Operator) – can do everything apart from adding/removing users and modifying their roles. This role is focused mainly on host and pool management (i.e.. creating storage, making pools, managing the hosts etc.)
- *Virtual Machine Power Administrator* (VM Power Admin) – creates and manages Virtual Machines. This role is focused on provisioning VMs for use by a VM operator.
- *Virtual Machine Administrator* (VM Admin) – similar to a VM Power Admin, but cannot migrate VMs or perform snapshots.
- *Virtual Machine Operator* (VM Operator) – similar to VM Admin, but cannot create/destroy VMs – but can perform start/stop lifecycle operations.
- *Read-only* (Read Only) – can view resource pool and performance data.

Note:

You cannot add, remove or modify roles in this version of XenServer.

Warning:

You can not assign the role of *pool-admin* to an AD group which has more than 500 members, if you want users of the AD group to have SSH access.

For a summary of the permissions available for each role and more detailed information on the operations available for each permission, see [the section called "Definitions of RBAC roles and permissions"](#).

All XenServer users need to be allocated to an appropriate role. By default, all new users will be allocated to the Pool Administrator role. It is possible for a user to be assigned to multiple roles; in that scenario, the user will have the union of all the permissions of all their assigned roles.

A user's role can be changed in two ways:

1. Modify the subject -> role mapping (this requires the assign/modify role permission, only available to a Pool Administrator.)
2. Modify the user's containing group membership in Active Directory.

Definitions of RBAC roles and permissions

The following table summarizes which permissions are available for each role. For details on the operations available for each permission, see Definitions of permissions.

Table 1. Permissions available for each role

Role permissions	Pool Admin	Pool Operator	VM Power Admin	VM Admin	VM Operator	Read Only
Assign/modify roles	X					
Log in to (physical) server consoles (through SSH and XenCenter)	X					
Server backup/restore	X					
Log out active user connections	X	X				
Create and dismiss alerts	X	X				
Cancel task of any user	X	X				
Pool management	X	X				
VM advanced operations	X	X	X			
VM create/destroy operations	X	X	X	X		

Role permissions	Pool Admin	Pool Operator	VM Power Admin	VM Admin	VM Operator	Read Only
VM change CD media	X	X	X	X	X	
View VM consoles	X	X	X	X	X	
XenCenter view mgmt ops	X	X	X	X	X	
Cancel own tasks	X	X	X	X	X	X
Read audit logs	X	X	X	X	X	X
Configure, Initialize, Enable, Disable WLB	X	X				
Apply WLB Optimization Recommendation	X	X				
Modify WLB Report Subscriptions	X	X				
Accept WLB Placement Recommendation	X	X	X			
Display WLB Configuration	X	X	X	X	X	X
Generate WLB Reports	X	X	X	X	X	
Connect to pool and read all pool metadata	X	X	X	X	X	X

Definitions of permissions

The following table provides additional details about permissions:

Table 2. Definitions of permissions

Permission	Allows Assignee To	Rationale/Comments
Assign/modify roles	<ul style="list-style-type: none"> Add/remove users Add/remove roles from users Enable and disable Active Directory integration (being joined to the domain) 	<p>This permission lets the user grant himself or herself any permission or perform any task.</p> <p>Warning: This role lets the user disable the Active Directory integration and all subjects added from Active Directory.</p>
Log in to server consoles	<ul style="list-style-type: none"> Server console access through ssh Server console access through XenCenter 	<p>Warning: With access to a root shell, the assignee could arbitrarily reconfigure the entire system, including RBAC.</p>
Server backup/restore VM create/destroy operations	<ul style="list-style-type: none"> Back up and restore servers Back up and restore pool metadata 	<p>The ability to restore a backup lets the assignee revert RBAC configuration changes.</p>
Log out active user connections	<ul style="list-style-type: none"> Ability to disconnect logged in users 	
Create/dismiss alerts		<p>Warning: A user with this permission can dismiss alerts for the entire pool.</p> <p>Note: The ability to view alerts is part of the Connect to Pool and read all pool metadata permission.</p>
Cancel task of any user	<ul style="list-style-type: none"> Cancel any user's running task 	<p>This permission lets the user request XenServer cancel an in-progress task initiated by any user.</p>

Permission	Allows Assignee To	Rationale/Comments
Pool management	<ul style="list-style-type: none"> • Set pool properties (naming, default SRs) • Enable, disable, and configure HA • Set per-VM HA restart priorities • Enable, disable, and configure Workload Balancing (WLB) • Add and remove server from pool • Emergency transition to master • Emergency master address • Emergency recover slaves • Designate new master • Manage pool and server certificates • Patching • Set server properties • Configure server logging • Enable and disable servers • Shut down, reboot, and power-on servers • System status reports • Apply license • Live migration of all other VMs on a server to another server, due to either WLB, Maintenance Mode, or HA • Configure server management interfaces • Disable server management • Delete crashdumps • Add, edit, and remove networks • Add, edit, and remove PBDs/PIFs/VLANs/Bonds/SRs • Add, remove, and retrieve secrets 	<p>This permission includes all the actions required to maintain a pool.</p> <p>Note: If the management interface is not functioning, no logins can authenticate except local root logins.</p>

Permission	Allows Assignee To	Rationale/Comments
VM advanced operations	<ul style="list-style-type: none"> Adjust VM memory (through Dynamic Memory Control) Create a VM snapshot with memory, take VM snapshots, and roll-back VMs Migrate VMs Start VMs, including specifying physical server Resume VMs 	This permission provides the assignee with enough privileges to start a VM on a different server if they are not satisfied with the server XenServer selected.
VM create/destroy operations	<ul style="list-style-type: none"> Install or delete Clone VMs Add, remove, and configure virtual disk/CD devices Add, remove, and configure virtual network devices Import/export VMs VM configuration change 	
VM change CD media	<ul style="list-style-type: none"> Eject current CD Insert new CD 	
VM change power state	<ul style="list-style-type: none"> Start VMs (automatic placement) Shut down VMs Reboot VMs Suspend VMs Resume VMs (automatic placement) 	This permission does not include start_on, resume_on, and migrate, which are part of the VM advanced operations permission.
View VM consoles	<ul style="list-style-type: none"> See and interact with VM consoles 	This permission does not let the user view server consoles.
Configure, Initialize, Enable, Disable WLB	<ul style="list-style-type: none"> Configure WLB Initialize WLB and change WLB servers Enable WLB Disable WLB 	When a user's role does not have this permission, this functionality is not visible.
Apply WLB Optimization Recommendations	<ul style="list-style-type: none"> Apply any optimization recommendations that appear in the WLB tab 	
Modify WLB Report Subscriptions	<ul style="list-style-type: none"> Change the WLB report generated or its recipient 	

Permission	Allows Assignee To	Rationale/Comments
Accept WLB Placement Recommendations	<ul style="list-style-type: none"> Select one of the servers Workload Balancing recommends for placement ("star" recommendations) 	
Display WLB Configuration	<ul style="list-style-type: none"> View WLB settings for a pool as shown on the WLB tab 	
Generate WLB Reports	<ul style="list-style-type: none"> View and run WLB reports, including the Pool Audit Trail report 	
XenCenter view mgmt operations	<ul style="list-style-type: none"> Create and modify global XenCenter folders Create and modify global XenCenter custom fields Create and modify global XenCenter searches 	Folders, custom fields, and searches are shared between all users accessing the pool
Cancel own tasks	<ul style="list-style-type: none"> Lets a user cancel their own tasks 	
Read audit log	<ul style="list-style-type: none"> Download the XenServer audit log 	
Connect to pool and read all pool metadata	<ul style="list-style-type: none"> Log in to pool View pool metadata View historical performance data View logged in users View users and roles View messages Register for and receive events 	

Note:

In some cases, a Read Only user cannot move a resource into a folder in XenCenter, even after receiving an elevation prompt and supplying the credentials of a more privileged user. In this case, log on to XenCenter as the more privileged user and retry the action.

Working with RBAC using the xe CLI

To list all the available defined roles in XenServer

- Run the command: `xe role-list`

This command returns a list of the currently defined roles, for example:



```
uuid( RO): 0165f154-ba3e-034e-6b27-5d271af109ba
name ( RO): pool-admin
description ( RO): The Pool Administrator role can do anything

uuid ( RO): b9ce9791-0604-50cd-0649-09b3284c7dfd
name ( RO): pool-operator
description ( RO): The Pool Operator can do anything but access Dom0 and mange subjects and roles

uuid( RO): 7955168d-7bec-10ed-105f-c6a7e6e63249
name ( RO): vm-power-admin
description ( RO): The VM Power Administrator role can do anything \
affecting VM properties across the pool

uuid ( RO): aaa00ab5-7340-bfbc-0dlb-7cf342639a6e
name ( RO): vm-admin
description ( RO): The VM Administrator role can do anything to a VM

uuid ( RO): fb8d4ff9-310c-a959-0613-54101535d3d5
name ( RO): vm-operator
description ( RO): The VM Operator role can do anything to an already

uuid ( RO): 7233b8e3-eacb-d7da-2c95-f2e581cdbf4e
name ( RO): read-only
description ( RO): The Read-Only role can only read values
```

Note:

The list of roles is static, so it is not possible to add or remove roles from it, only to list the available static ones.

To display a list of current subjects:

- Run the command `xe subject-list`

This will return a list of XenServer users, their uuid, and the roles they are associated with:



```
uuid ( RO): bb6dd239-1fa9-a06b-a497-3be28b8dca44
subject-identifier ( RO): S-1-5-21-1539997073-1618981536-2562117463-2244
other-config (MRO): subject-name: example01\user_vm_admin; subject-upn: \
  user_vm_admin@XENDT.NET; subject-uid: 1823475908; subject-gid: 1823474177; \
  subject-sid: S-1-5-21-1539997073-1618981536-2562117463-2244; subject-gecos: \
  user_vm_admin; subject-displayname: user_vm_admin; subject-is-group: false; \
  subject-account-disabled: false; subject-account-expired: false; \
  subject-account-locked: false; subject-password-expired: false
roles (SRO): vm-admin

uuid ( RO): 4fe89a50-6ala-d9dd-afb9-b554cd00c01a
subject-identifier ( RO): S-1-5-21-1539997073-1618981536-2562117463-2245
other-config (MRO): subject-name: example02\user_vm_op; subject-upn: \
  user_vm_op@XENDT.NET; subject-uid: 1823475909; subject-gid: 1823474177; \
  subject-sid: S-1-5-21-1539997073-1618981536-2562117463-2245; \
  subject-gecos: user_vm_op; subject-displayname: user_vm_op; \
  subject-is-group: false; subject-account-disabled: false; \
  subject-account-expired: false; subject-account-locked: \
  false; subject-password-expired: false
roles (SRO): vm-operator

uuid ( RO): 8a63fbf0-9ef4-4fef-b4a5-b42984c27267
subject-identifier ( RO): S-1-5-21-1539997073-1618981536-2562117463-2242
other-config (MRO): subject-name: example03\user_pool_op; \
  subject-upn: user_pool_op@XENDT.NET; subject-uid: 1823475906; \
  subject-gid: 1823474177; subject-s id:
  S-1-5-21-1539997073-1618981536-2562117463-2242; \
  subject-gecos: user_pool_op; subject-displayname: user_pool_op; \
  subject-is-group: false; subject-account-disabled: false; \
  subject-account-expired: false; subject-account-locked: \
  false; subject-password-expired: false
roles (SRO): pool-operator
```

To add a subject to RBAC

In order to enable existing AD users to use RBAC, you will need to create a subject instance within XenServer, either for the AD user directly, or for one of their containing groups:

1. Run the command `xe subject-add subject-name=<AD user/group>`

This adds a new subject instance.

To assign an RBAC role to a created subject

Once you have added a subject, you can assign it to an RBAC role. You can refer to the role by either its uuid or name:

1. Run the command:

```
xe subject-role-add uuid=<subject uuid> role-uuid=<role uuid>
```

or

```
xe subject-role-add uuid=<subject uuid> role-name=<role name>
```

For example, the following command adds a subject with the uuid b9b3d03b-3d10-79d3-8ed7-a782c5ea13b4 to the Pool Administrator role:

```
xe subject-role-add uuid=b9b3d03b-3d10-79d3-8ed7-a782c5ea13b4 role-name=pool-admin
```



To change a subject's RBAC role:

To change a user's role it is necessary to remove them from their existing role, and add them to a new role:

1. Run the commands:

```
xe subject-role-remove uuid=<subject uuid> role-name= \  
    <role_name_to_remove>  
xe subject-role-add uuid=<subject uuid > role-name= \  
    <role_name_to_add>
```

To ensure that the new role takes effect, the user should be logged out and logged back in again (this requires the "Logout Active User Connections" permission - available to a Pool Administrator or Pool Operator).

Warning:

Once you have added or removed a pool-admin subject, there can be a delay of a few seconds for ssh sessions associated to this subject to be accepted by all hosts of the pool.

Auditing

The RBAC audit log will record any operation taken by a logged-in user.

- the message will explicitly record the Subject ID and user name associated with the session that invoked the operation.
- if an operation is invoked for which the subject does not have authorization, this will be logged.
- if the operation succeeded then this is recorded; if the operation failed then the error code is logged.

Audit log xe CLI commands

```
xe audit-log-get [since=<timestamp>] filename=<output filename>
```

This command downloads to a file all the available records of the RBAC audit file in the pool. If the optional parameter 'since' is present, then it only downloads the records from that specific point in time.

To obtain all audit records from the pool

Run the following command:

```
xe audit-log-get filename=/tmp/auditlog-pool-actions.out
```

To obtain audit records of the pool since a precise millisecond timestamp

Run the following command:

```
xe audit-log-get since=2009-09-24T17:56:20.530Z \  
filename=/tmp/auditlog-pool-actions.out
```

To obtain audit records of the pool since a precise minute timestamp

Run the following command:

```
xe audit-log-get since=2009-09-24T17:56Z \  
filename=/tmp/auditlog-pool-actions.out
```



How does XenServer compute the roles for the session?

1. The subject is authenticated via the Active Directory server to verify which containing groups the subject may also belong to.
2. XenServer then verifies which roles have been assigned both to the subject, and to its containing groups.
3. As subjects can be members of multiple Active Directory groups, they will inherit all of the permissions of the associated roles.

In this illustration, since Subject 2 (Group 2) is the Pool Operator and User 1 is a member of Group 2, when Subject 3 (User 1) tries to log in, he or she inherits both Subject 3 (VM Operator) and Group 2 (Pool Operator) roles. Since the Pool Operator role is higher, the resulting role for Subject 3 (User 1) is Pool Operator and not VM Operator.

XenServer hosts and resource pools

This chapter describes how resource pools can be created through a series of examples using the `xe` command line interface (CLI). A simple NFS-based shared storage configuration is presented and a number of simple VM management examples are discussed. Procedures for dealing with physical node failures are also described.

Hosts and resource pools overview

A *resource pool* comprises multiple XenServer host installations, bound together into a single managed entity which can host Virtual Machines. When combined with shared storage, a resource pool enables VMs to be started on *any* XenServer host which has sufficient memory and then dynamically moved between XenServer hosts while running with minimal downtime (XenMotion). If an individual XenServer host suffers a hardware failure, then the administrator can restart the failed VMs on another XenServer host in the same resource pool. If high availability (HA) is enabled on the resource pool, VMs will automatically be moved if their host fails. Up to 16 hosts are supported per resource pool, although this restriction is not enforced.

A pool always has at least one physical node, known as the *master*. Only the master node exposes an administration interface (used by XenCenter and the XenServer Command Line Interface, known as the `xe` CLI); the master forwards commands to individual members as necessary.

Note:

If the pool's master fails, master re-election will only take place if High Availability is enabled.

Requirements for creating resource pools

A resource pool is a homogeneous (or heterogeneous with restrictions, see [the section called "Creating heterogeneous resource pools"](#)) aggregate of one or more XenServer hosts, up to a maximum of 16. The definition of homogeneous is:

- the CPUs on the server joining the pool are the same (in terms of vendor, model, and features) as the CPUs on servers already in the pool.
- the server joining the pool is running the same version of XenServer software, at the same patch level, as servers already in the pool

The software will enforce additional constraints when joining a server to a pool – in particular:

- it is not a member of an existing resource pool
- it has no shared storage configured
- there are no running or suspended VMs on the XenServer host which is joining
- there are no active operations on the VMs in progress, such as one shutting down

You must also check that the clock of the host joining the pool is synchronized to the same time as the pool master (for example, by using NTP), that its management interface is not bonded (you can configure this once the host has successfully joined the pool), and that its management IP address is static (either configured on the host itself or by using an appropriate configuration on your DHCP server).

XenServer hosts in resource pools may contain different numbers of physical network interfaces and have local storage repositories of varying size. In practice, it is often difficult to obtain multiple servers with the exact same CPUs, and so minor variations are permitted. If you are sure that it is acceptable in your environment for hosts with varying CPUs to be part of the same resource pool, then the pool joining operation can be forced by passing a `--force` parameter.

Note:

The requirement for a XenServer host to have a static IP address to be part of a resource pool also applies to servers providing shared NFS or iSCSI storage for the pool.

Although not a strict technical requirement for creating a resource pool, the advantages of pools (for example, the ability to dynamically choose on which XenServer host to run a VM and to dynamically move a VM between XenServer hosts) are only available if the pool has one or more shared storage repositories. If possible, postpone creating a pool of XenServer hosts until shared storage is available. Once shared storage has been added, Citrix recommends that you move existing VMs whose disks are in local storage into shared storage. This can be done using the **xe vm-copy** command or XenCenter.

Creating a resource pool

Resource pools can be created using either the XenCenter management console or the CLI. When you join a new host to a resource pool, the joining host synchronizes its local database with the pool-wide one, and inherits some settings from the pool:

- VM, local, and remote storage configuration is added to the pool-wide database. All of these will still be tied to the joining host in the pool unless you explicitly take action to make the resources shared after the join has completed.
- The joining host inherits existing shared storage repositories in the pool and appropriate PBD records are created so that the new host can access existing shared storage automatically.
- Networking information is partially inherited to the joining host: the *structural* details of NICs, VLANs and bonded interfaces are all inherited, but *policy* information is not. This policy information, which must be re-configured, includes:
 - the IP addresses of management NICs, which are preserved from the original configuration
 - the location of the management interface, which remains the same as the original configuration. For example, if the other pool hosts have their management interface on a bonded interface, then the joining host must be explicitly migrated to the bond once it has joined. See [To add NIC bonds to the pool master and other hosts](#) for details on how to migrate the management interface to a bond.
 - Dedicated storage NICs, which must be re-assigned to the joining host from XenCenter or the CLI, and the PBDs re-plugged to route the traffic accordingly. This is because IP addresses are not assigned as part of the pool join operation, and the storage NIC is not useful without this configured correctly. See [the section called “Configuring a dedicated storage NIC”](#) for details on how to dedicate a storage NIC from the CLI.

To join XenServer hosts *host1* and *host2* into a resource pool using the CLI

1. Open a console on XenServer host *host2*.
2. Command XenServer host *host2* to join the pool on XenServer host *host1* by issuing the command:

```
xe pool-join master-address=<host1> master-username=<administrators_username> \
master-password=<password>
```

The *master-address* must be set to the fully-qualified domain name of XenServer host *host1* and the *password* must be the administrator password set when XenServer host *host1* was installed.

Naming a resource pool

- XenServer hosts belong to an unnamed pool by default. To create your first resource pool, rename the existing nameless pool. Use tab-complete to find the `pool_uuid`:



```
xe pool-param-set name-label=<"New Pool"> uuid=<pool_uuid>
```

Creating heterogeneous resource pools

Note:

Heterogeneous resource pool creation is only available for XenServer Enterprise or Platinum editions. To learn more about XenServer editions and to find out how to upgrade, visit the Citrix website [here](#)

XenServer 5.6 simplifies expanding deployments over time by allowing disparate host hardware to be joined into a resource pool, known as heterogeneous resource pools. Heterogeneous resource pools are made possible by leveraging technologies in recent Intel (FlexMigration) and AMD (Extended Migration) CPUs that provide CPU "masking" or "leveling". These features allow a CPU to be configured to *appear* as providing a different make, model, or functionality than it actually does. This enables you to create pools of hosts with disparate CPUs but still safely support live migrations.

Using XenServer to mask the CPU features of a new server, so that it will match the features of the existing servers in a pool, requires the following:

- the CPUs of the server joining the pool must be of the same vendor (i.e.. AMD, Intel) as the CPUs on servers already in the pool, though the specific type, (family, model and stepping numbers) need not be.
- the CPUs of the server joining the pool must support either Intel FlexMigration or AMD Enhanced Migration.
- the features of the older CPUs must be a sub-set of the features of the CPUs of the server joining the pool.
- the server joining the pool is running the same version of XenServer software, with the same hotfixes installed, as servers already in the pool.
- an Enterprise or Platinum license.

Creating heterogeneous resource pools is most easily done with XenCenter which will automatically suggest using CPU masking when possible. Refer to the *Pool Requirements* section in the XenCenter help for more details. To display the help in XenCenter press F1.

To add a heterogeneous XenServer host to a resource pool using the xe CLI

1. Find the CPU features of the Pool Master by running the **xe host-get-cpu-features** command.
2. On the new server, run the **xe host-set-cpu-features** command and copy and paste the Pool Master's features into the `features` parameter. For example:

```
xe host-set-cpu-features features=<pool_master's_cpu_features>
```

3. Restart the new server.
4. Run the **xe pool-join** command on the new server to join the pool.

To return a server with masked CPU features back to its normal capabilities, run the **xe host-reset-cpu-features** command.

Note:

To display a list of all properties of the CPUs in a host, run the **xe host-cpu-info** command.

Adding shared storage

For a complete list of supported shared storage types, see the [Storage chapter](#). This section demonstrates how shared storage (represented as a storage repository) can be created on an existing NFS server.

Adding NFS shared storage to a resource pool using the CLI

1. Open a console on any XenServer host in the pool.
2. Create the storage repository on `<server:/path>` by issuing the command

```
xe sr-create content-type=user type=nfs name-label=<"Example SR"> shared=true \
  device-config:server=<server> \
  device-config:serverpath=<path>
```

The `device-config:server` refers to the hostname of the NFS server and `device-config:serverpath` refers to the path on the NFS server. Since `shared` is set to true, the shared storage will be automatically connected to every XenServer host in the pool and any XenServer hosts that subsequently join will also be connected to the storage. The Universally Unique Identifier (UUID) of the created storage repository will be printed on the screen.

3. Find the UUID of the pool by the command

```
xe pool-list
```

4. Set the shared storage as the pool-wide default with the command

```
xe pool-param-set uuid=<pool_uuid> default-SR=<sr_uuid>
```

Since the shared storage has been set as the pool-wide default, all future VMs will have their disks created on shared storage by default. See [Storage](#) for information about creating other types of shared storage.

Removing a XenServer host from a resource pool

When a XenServer host is removed (*ejected*) from a pool, the machine is rebooted, reinitialized, and left in a state equivalent to that after a fresh installation. It is important not to eject a XenServer host from a pool if there is important data on the local disks.

To remove a host from a resource pool using the CLI

1. Open a console on any host in the pool.
2. Find the UUID of the host by running the command

```
xe host-list
```

3. Eject the required host from the pool:

```
xe pool-eject host-uuid=<host_uuid>
```

The XenServer host will be ejected and left in a freshly-installed state.

Warning:

Do *not* eject a host from a resource pool if it contains important data stored on its local disks. All of the data will be erased upon ejection from the pool. If you wish to preserve this data, copy the VM to shared storage on the pool first using XenCenter, or the **xe vm-copy** CLI command.

When a XenServer host containing locally stored VMs is ejected from a pool, those VMs will still be present in the pool database and visible to the other XenServer hosts. They will not start until the virtual disks associated with them have been changed to point at shared storage which can be seen by other XenServer hosts in the pool, or simply removed. It is for this reason that you are strongly advised to move any local storage to shared storage upon joining a pool, so that individual XenServer hosts can be ejected (or physically fail) without loss of data.



High Availability

This section explains the XenServer implementation of virtual machine high availability (HA), and how to configure it using the xe CLI.

Note:

XenServer HA is only available with XenServer Advanced edition or above. To find out about XenServer editions, visit the Citrix website [here](#).

HA Overview

When HA is enabled, XenServer continually monitors the health of the hosts in a pool. The HA mechanism automatically moves protected VMs to a healthy host if the current VM host fails. Additionally, if the host that fails is the master, HA selects another host to take over the master role automatically, so that you can continue to manage the XenServer pool.

To absolutely guarantee that a host is unreachable, a resource pool configured for high-availability uses several *heartbeat* mechanisms to regularly check up on hosts. These heartbeats go through both the storage interfaces (to the *Heartbeat SR*) and the networking interfaces (over the management interfaces). Both of these heartbeat routes can be multi-homed for additional resilience to prevent false positives.

XenServer dynamically maintains a *failover plan* which details what to do if a set of hosts in a pool fail at any given time. An important concept to understand is the *host failures to tolerate* value, which is defined as part of HA configuration. This determines the number of failures that is allowed without any loss of service. For example, if a resource pool consisted of 16 hosts, and the tolerated failures is set to 3, the pool calculates a failover plan that allows for any 3 hosts to fail and still be able to restart VMs on other hosts. If a plan cannot be found, then the pool is considered to be *overcommitted*. The plan is dynamically recalculated based on VM lifecycle operations and movement. Alerts are sent (either through XenCenter or e-mail) if changes (for example the addition on new VMs to the pool) cause your pool to become overcommitted.

Overcommitting

A pool is overcommitted if the VMs that are currently running could not be restarted elsewhere following a user-defined number of host failures.

This would happen if there was not enough free memory across the pool to run those VMs following failure. However there are also more subtle changes which can make HA guarantees unsustainable: changes to Virtual Block Devices (VBDs) and networks can affect which VMs may be restarted on which hosts. Currently it is not possible for XenServer to check all actions before they occur and determine if they will cause violation of HA demands. However an asynchronous notification is sent if HA becomes unsustainable.

Overcommitment Warning

If you attempt to start or resume a VM and that action causes the pool to be overcommitted, a warning alert is raised. This warning is displayed in XenCenter and is also available as a message instance through the Xen API. The message may also be sent to an email address if configured. You will then be allowed to cancel the operation, or proceed anyway. Proceeding causes the pool to become overcommitted. The amount of memory used by VMs of different priorities is displayed at the pool and host levels.

Host Fencing

If a server failure occurs such as the loss of network connectivity or a problem with the control stack is encountered, the XenServer host self-fences to ensure that the VMs are not running on two servers simultaneously. When a fence action is taken, the server immediately and abruptly restarts, causing all VMs



running on it to be stopped. The other servers will detect that the VMs are no longer running and the VMs will be restarted according to the restart priorities assigned to them. The fenced server will enter a reboot sequence, and when it has restarted it will try to re-join the resource pool.

Configuration Requirements

To use the HA feature, you need:

- Shared storage, including at least one iSCSI or Fibre Channel LUN of size 356MB or greater - the *heartbeat* SR. The HA mechanism creates two volumes on the heartbeat SR:

4MB heartbeat volume
Used for heartbeating.

256MB metadata volume
Stores pool master metadata to be used in the case of master failover.

If you are using a NetApp or EqualLogic SR, manually provision an iSCSI LUN on the array to use as the heartbeat SR.

- A XenServer pool (this feature provides high availability at the server level within a single resource pool).
- Enterprise licenses on all hosts.
- Static IP addresses for all hosts.

Warning:

Should the IP address of a server change while HA is enabled, HA will assume that the host's network has failed, and will probably fence the host and leave it in an unbootable state. To remedy this situation, disable HA using the **host-emergency-ha-disable** command, reset the pool master using **pool-emergency-reset-master**, and then re-enable HA.

For a VM to be protected by the HA feature, it must be agile. This means that:

- it must have its virtual disks on shared storage (any type of shared storage may be used; the iSCSI or Fibre Channel LUN is only required for the storage heartbeat and can be used for virtual disk storage if you prefer, but this is not necessary)
- it must not have a connection to a local DVD drive configured
- it should have its virtual network interfaces on pool-wide networks.

Citrix strongly recommends the use of a bonded management interface on the servers in the pool if HA is enabled, and multipathed storage for the heartbeat SR.

If you create VLANs and bonded interfaces from the CLI, then they may not be plugged in and active despite being created. In this situation, a VM can appear to be not agile, and cannot be protected by HA. If this occurs, use the CLI **pif-plug** command to bring the VLAN and bond PIFs up so that the VM can become agile. You can also determine precisely why a VM is not agile by using the **xe diagnostic-vm-status** CLI command to analyze its placement constraints, and take remedial action if required.

Restart priorities

Virtual machines are assigned a restart priority and a flag that indicates whether they should be protected by HA or not. When HA is enabled, every effort is made to keep protected virtual machines live. If a restart priority is specified, any protected VM that is halted will be started automatically. If a server fails then the VMs on it will be started on another server.

The possible restart priorities are:

1 | 2 | 3

when a pool is overcommitted the HA mechanism will attempt to restart protected VMs with the lowest restart priority first

best-effort

VMs with this priority setting will be restarted only when the system has attempted to restart protected VMs

ha-always-run=false

VMs with this parameter set will not be restarted

The restart priorities determine the order in which VMs are restarted when a failure occurs. In a given configuration where a number of server failures greater than zero can be tolerated (as indicated in the HA panel in the GUI, or by the *ha-plan-exists-for* field on the pool object on the CLI), the VMs that have restart priorities 1, 2 or 3 are guaranteed to be restarted given the stated number of server failures. VMs with a *best-effort* priority setting are not part of the failover plan and are not guaranteed to be kept running, since capacity is not reserved for them. If the pool experiences server failures and enters a state where the number of tolerable failures drops to zero, the protected VMs will no longer be guaranteed to be restarted. If this condition is reached, a system alert will be generated. In this case, should an additional failure occur, all VMs that have a restart priority set will behave according to the *best-effort* behavior.

If a protected VM cannot be restarted at the time of a server failure (for example, if the pool was overcommitted when the failure occurred), further attempts to start this VM will be made as the state of the pool changes. This means that if extra capacity becomes available in a pool (if you shut down a non-essential VM, or add an additional server, for example), a fresh attempt to restart the protected VMs will be made, which may now succeed.

Note:

No running VM will ever be stopped or migrated in order to free resources for a VM with *always-run=true* to be restarted.

Enabling HA on a XenServer pool

HA can be enabled on a pool using either XenCenter or the command-line interface. In either case, you will specify a set of priorities that determine which VMs should be given highest restart priority when a pool is overcommitted.

Warning:

When HA is enabled, some operations that would compromise the plan for restarting VMs may be disabled, such as removing a server from a pool. To perform these operations, HA can be temporarily disabled, or alternately, VMs protected by HA made unprotected.

Enabling HA using the CLI

1. Verify that you have a compatible Storage Repository (SR) attached to your pool. iSCSI or Fibre Channel are compatible SR types. Please refer to the reference guide for details on how to configure such a storage repository using the CLI.

2. For each VM you wish to protect, set a restart priority. You can do this as follows:

```
xe vm-param-set uuid=<vm_uuid> ha-restart-priority=<1> ha-always-run=true
```

3. Enable HA on the pool:

```
xe pool-ha-enable heartbeat-sr-uuids=<sr_uuid>
```

4. Run the **pool-ha-compute-max-host-failures-to-tolerate** command. This command returns the maximum number of hosts that can fail before there are insufficient resources to run all the protected VMs in the pool.

```
xe pool-ha-compute-max-host-failures-to-tolerate
```

The number of failures to tolerate determines when an alert is sent: the system will recompute a failover plan as the state of the pool changes and with this computation the system identifies the capacity of the pool and how many more failures are possible without loss of the liveness guarantee for protected VMs. A system alert is generated when this computed value falls below the specified value for *ha-host-failures-to-tolerate*.

5. Specify the number of failures to tolerate parameter. This should be less than or equal to the computed value:

```
xe pool-param-set ha-host-failures-to-tolerate=<2>
```

Removing HA protection from a VM using the CLI

To disable HA features for a VM, use the **xe vm-param-set** command to set the *ha-always-run* parameter to *false*. This does not clear the VM restart priority settings. You can enable HA for a VM again by setting the *ha-always-run* parameter to *true*.

Recovering an unreachable host

If for some reason a host cannot access the HA statefile, it is possible that a host may become unreachable. To recover your XenServer installation it may be necessary to disable HA using the **host-emergency-ha-disable** command:

```
xe host-emergency-ha-disable --force
```

If the host was the pool master, then it should start up as normal with HA disabled. Slaves should reconnect and automatically disable HA. If the host was a Pool slave and cannot contact the master, then it may be necessary to force the host to reboot as a pool master (**xe pool-emergency-transition-to-master**) or to tell it where the new master is (**xe pool-emergency-reset-master**):

```
xe pool-emergency-transition-to-master uuid=<host_uuid>
xe pool-emergency-reset-master master-address=<new_master_hostname>
```

When all hosts have successfully restarted, re-enable HA:

```
xe pool-ha-enable heartbeat-sr-uuid=<sr_uuid>
```

Shutting down a host when HA is enabled

When HA is enabled special care needs to be taken when shutting down or rebooting a host to prevent the HA mechanism from assuming that the host has failed. To shutdown a host cleanly in an HA-enabled environment, first *disable* the host, then *evacuate* the host and finally *shutdown* the host using either XenCenter or the CLI. To shutdown a host in an HA-enabled environment on the command line:

```
xe host-disable host=<host_name>
xe host-evacuate uuid=<host_uuid>
xe host-shutdown host=<host_name>
```

Shutting down a VM when it is protected by HA

When a VM is protected under a HA plan and set to restart automatically, it cannot be shut down while this protection is active. To shut down a VM, first disable its HA protection and then execute the CLI command.



XenCenter offers you a dialog box to automate disabling the protection if you click on the **Shutdown** button of a protected VM.

Note:

If you shut down a VM from within the guest, and the VM is protected, it is automatically restarted under the HA failure conditions. This helps ensure that operator error (or an errant program that mistakenly shuts down the VM) does not result in a protected VM being left shut down accidentally. If you want to shut this VM down, disable its HA protection first.

Host Power On

Powering on hosts remotely

You can use the XenServer Host Power On feature to turn a server on and off remotely, either from XenCenter or by using the CLI. When using Workload Balancing (WLB), you can configure Workload Balancing to turn hosts on and off automatically as VMs are consolidated or brought back online.

To enable host power, the server must have one of the following power-control solutions:

- **Wake On LAN enabled network card.**
- **Dell Remote Access Cards (DRAC).** To use XenServer with DRAC, you must install the Dell supplemental pack to get DRAC support. DRAC support requires installing RACADM command-line utility on the server with the remote access controller and enable DRAC and its interface. RACADM is often included in the DRAC management software. For more information, see Dell's DRAC documentation.
- **Hewlett-Packard Integrated Lights-Out (iLO).** To use XenServer with iLO, you must enable iLO on the host and connect interface to the network. For more information, see HP's iLO documentation.
- A custom script based on the XenAPI that enables you to turn the power on and off through XenServer. For more information, see [Configuring a Custom Script for XenServer's Host Power On Feature].

Using the Host Power On feature requires three tasks:

1. Ensuring the hosts in the pool support controlling the power remotely (that is, they have Wake-on-LAN functionality, a DRAC or iLO card, or you created custom script).
2. Enabling the Host Power On functionality using the CLI or XenCenter.
3. (Optional.) Configuring automatic Host Power On functionality in Workload Balancing. See [the section called "Optimizing and Managing Power Automatically"](#).

Note:

You must enable Host Power On and configure the Power Management feature in Workload Balancing before Workload Balancing can turn hosts on and off automatically.

Using the CLI to Manage Host Power On

You can manage the Host Power On feature using either the CLI or XenCenter. This topic provides information about managing it with the CLI.

Host Power On is enabled at the host level (that is, on each XenServer).

After you enable Host Power On, you can turn hosts on using either the CLI or XenCenter.

After configuration, you can configure and run the Workload Balancing Automation and Host Power Management features, as described in the Workload Balancing chapter. To enable Workload



Balancing's Host Power Management feature, use the **pool-send-wlb-configuration** command with the `ParticipatesInPowerManagement=<true>` \ `config:set_host_configuration=<true>` arguments.

To enable Host Power On using the CLI

1. Run the command:

```
xe host-set-power-on host=<host uuid>\
power-on-mode=( " " , "wake-on-lan",
"iLO", "DRAC", "custom")
power-on-config:key=value
```

For iLO and DRAC the keys are `power_on_ip`, `power_on_user`, `power_on_password`. Use `power_on_password` to specify the password if you are using the secret feature.

To turn on hosts remotely using the CLI

1. Run the command:

```
xe host-power-on host=<host uuid>
```

Configuring a Custom Script for XenServer's Host Power On Feature

If your servers' remote-power solution uses a protocol that is not supported by default (such as Wake-On-Ring or Intel Active Management Technology), you can create a custom Linux Python script to turn on your XenServer computers remotely. However, you can also create custom scripts for iLO, DRAC, and Wake-On-LAN remote-power solutions.

This topic provides information about configuring a custom script for Host Power On using the key/value pairs associated with the XenServer API call `host.power_on`.

When you create a custom script, run it from the command line each time you want to control power remotely on XenServer. Alternatively, you can specify it in XenCenter and use the XenCenter UI features to interact with it.

The XenServer API is documented in the document, the [Citrix XenServer Management API], which is available from the Citrix Web site.

Note:

Do not modify the scripts provided by default in the `/etc/xapi.d/plugins/` directory. You can include new scripts in this directory, but you should never modify the scripts contained in that directory after installation.

Key/Value Pairs

To use Host Power On, you must configure the `host.power_on_mode` and `host.power_on_config` keys. Their values are provided below.

There is also an API call that lets you set these fields all at once:

```
void host.set_host_power_on_mode(string mode, Dictionary<string,string> config)
```

`host.power_on_mode`

- **Definition:** This contains key/value pairs to specify the type of remote-power solution (for example, Dell DRAC).

- **Possible values:**
 - An empty string, representing power-control disabled
 - "iLO". Lets you specify HP iLO.
 - "DRAC". Lets you specify Dell DRAC. To use DRAC, you must have already installed the Dell supplemental pack.
 - "wake-on-lan". Lets you specify Wake on LAN.
 - Any other name (used to specify a custom power-on script). This option is used to specify a custom script for power management.
- **Type:** string

host.power_on_config

- **Definition:** This contains key/value pairs for mode configuration. Provides additional information for
- **Possible values:**
 - If you configured iLO or DRAC as the type of remote-power solution, you must also specify one of the following keys:
 - "power_on_ip". This is the IP address you specified configured to communicate with the power-control card. Alternatively, you can enter the domain name for the network interface where iLO or DRAC is configured.
 - "power_on_user". This is the iLO or DRAC user name that is associated with the management processor, which you may or may not have changed from its factory default settings.
 - "power_on_password_secret". Specifies using the secrets feature to secure your password.
 - To use the secrets feature to store your password, specify the key "power_on_password_secret".
- **Type:** Map (string,string)

Sample Script

This sample script imports the XenServer API, defines itself as a custom script, and then passes parameters specific to the host you want to control remotely. You must define the parameters *session*, *remote_host*, and *power_on_config* in all custom scripts.

The result is only displayed when the script is unsuccessful.

```
import XenAPI
def custom(session,remote_host,
power_on_config):
result="Power On Not Successful"
for key in power_on_config.keys():
result=result+
key="+key+"
value="+power_on_config[key]
return result
```

Note:

After creation, save the script in the /etc/xapi.d/plugins with a .py extension.

Storage

This chapter discusses the framework for storage abstractions. It describes the way physical storage hardware of various kinds is mapped to VMs, and the software objects used by the XenServer host API to perform storage-related tasks. Detailed sections on each of the supported storage types include procedures for creating storage for VMs using the CLI, with type-specific device configuration options, generating snapshots for backup purposes and some best practices for managing storage in XenServer host environments. Finally, the virtual disk QoS (quality of service) settings are described.

Storage Overview

This section explains what the XenServer storage objects are and how they are related to each other.

Storage Repositories (SRs)

XenServer defines a container called a storage repository (SR) to describe a particular storage target, in which Virtual Disk Images (VDIs) are stored. A VDI is a disk abstraction which contains the contents of a virtual disk.

The interface to storage hardware allows VDIs to be supported on a large number of SR types. The XenServer SR is very flexible, with built-in support for IDE, SATA, SCSI and SAS drives locally connected, and iSCSI, NFS, SAS and Fibre Channel remotely connected. The SR and VDI abstractions allow advanced storage features such as sparse provisioning, VDI snapshots, and fast cloning to be exposed on storage targets that support them. For storage subsystems that do not inherently support advanced operations directly, a software stack is provided based on Microsoft's Virtual Hard Disk (VHD) specification which implements these features.

Each XenServer host can use multiple SRs and different SR types simultaneously. These SRs can be shared between hosts or dedicated to particular hosts. Shared storage is pooled between multiple hosts within a defined resource pool. A shared SR must be network accessible to each host. All hosts in a single resource pool must have at least one shared SR in common.

SRs are storage targets containing virtual disk images (VDIs). SR commands provide operations for creating, destroying, resizing, cloning, connecting and discovering the individual VDIs that they contain.

A storage repository is a persistent, on-disk data structure. For SR types that use an underlying block device, the process of creating a new SR involves erasing any existing data on the specified storage target. Other storage types such as NFS, Netapp, Equallogic and StorageLink SRs, create a new container on the storage array in parallel to existing SRs.

CLI operations to manage storage repositories are described in [the section called “SR commands”](#).

Virtual Disk Images (VDIs)

Virtual Disk Images are a storage abstraction that is presented to a VM. VDIs are the fundamental unit of virtualized storage in XenServer. Similar to SRs, VDIs are persistent, on-disk objects that exist independently of XenServer hosts. CLI operations to manage VDIs are described in [the section called “VDI commands”](#). The actual on-disk representation of the data differs by the SR type and is managed by a separate storage plugin interface for each SR, called the SM API.

Physical Block Devices (PBDs)

Physical Block Devices represent the interface between a physical server and an attached SR. PBDs are connector objects that allow a given SR to be mapped to a XenServer host. PBDs store the device

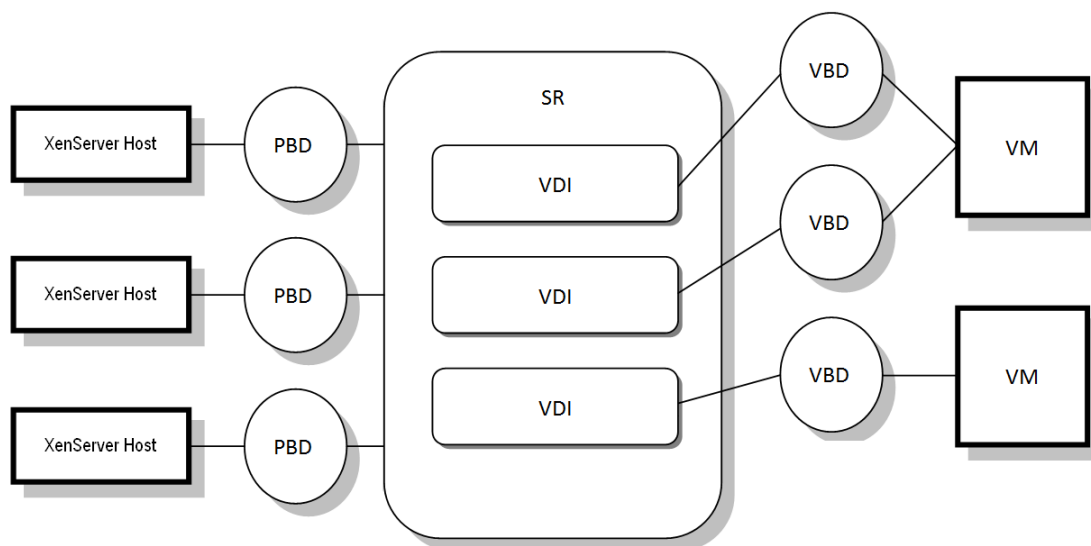
configuration fields that are used to connect to and interact with a given storage target. For example, NFS device configuration includes the IP address of the NFS server and the associated path that the XenServer host mounts. PBD objects manage the run-time attachment of a given SR to a given XenServer host. CLI operations relating to PBDs are described in [the section called “PBD commands”](#).

Virtual Block Devices (VBDs)

Virtual Block Devices are connector objects (similar to the PBD described above) that allows mappings between VDIs and VMs. In addition to providing a mechanism for attaching (also called plugging) a VDI into a VM, VBDs allow for the fine-tuning of parameters regarding QoS (quality of service), statistics, and the bootability of a given VDI. CLI operations relating to VBDs are described in [the section called “VBD commands”](#).

Summary of Storage objects

The following image is a summary of how the storage objects presented so far are related:



Graphical overview of storage repositories and related objects

Virtual Disk Data Formats

In general, there are three types of mapping of physical storage to a VDI:

- *File-based VHD on a Filesystem*; VM images are stored as thin-provisioned VHD format files on either a local non-shared Filesystem (EXT type SR) or a shared NFS target (NFS type SR)
- *Logical Volume-based VHD on a LUN*; The default XenServer blockdevice-based storage inserts a Logical Volume manager on a disk, either a locally attached device (LVM type SR) or a SAN attached LUN over either Fibre Channel (LVMoHBA type SR), iSCSI (LVMoISCSI type SR) or SAS (LVMoHBA type Sr). VDIs are represented as volumes within the Volume manager and stored in VHD format to allow thin provisioning of reference nodes on snapshot and clone.
- *LUN per VDI*; LUNs are directly mapped to VMs as VDIs by SR types that provide an array-specific plugin (Netapp, Equallogic or StorageLink type SRs). The array storage abstraction therefore matches the VDI storage abstraction for environments that manage storage provisioning at an array level.

VHD-based VDIs

VHD files may be *chained*, allowing two VDIs to share common data. In cases where a VHD-backed VM is cloned, the resulting VMs share the common on-disk data at the time of cloning. Each proceeds to make its own changes in an isolated copy-on-write (CoW) version of the VDI. This feature allows VHD-based VMs to be quickly cloned from templates, facilitating very fast provisioning and deployment of new VMs.

The VHD format used by LVM-based and File-based SR types in XenServer uses sparse provisioning. The image file is automatically extended in 2MB chunks as the VM writes data into the disk. For File-based VHD, this has the considerable benefit that VM image files take up only as much space on the physical storage as required. With LVM-based VHD the underlying logical volume container must be sized to the virtual size of the VDI, however unused space on the underlying CoW instance disk is reclaimed when a snapshot or clone occurs. The difference between the two behaviors can be characterized in the following way:

- For *LVM-based VHDs*, the difference disk nodes within the chain consume only as much data as has been written to disk but the leaf nodes (VDI clones) remain fully inflated to the virtual size of the disk. Snapshot leaf nodes (VDI snapshots) remain deflated when not in use and can be attached Read-only to preserve the deflated allocation. Snapshot nodes that are attached Read-Write will be fully inflated on attach, and deflated on detach.
- For *file-based VHDs*, all nodes consume only as much data as has been written, and the leaf node files grow to accommodate data as it is actively written. If a 100GB VDI is allocated for a new VM and an OS is installed, the VDI file will physically be only the size of the OS data that has been written to the disk, plus some minor metadata overhead.

When cloning VMs based off a single VHD template, each child VM forms a chain where new changes are written to the new VM, and old blocks are directly read from the parent template. If the new VM was converted into a further template and more VMs cloned, then the resulting chain will result in degraded performance. XenServer supports a maximum chain length of 30, but it is generally not recommended that you approach this limit without good reason. If in doubt, you can always "copy" the VM using XenServer or the **vm-copy** command, which resets the chain length back to 0.

VHD Chain Coalescing

VHD images support chaining, which is the process whereby information shared between one or more VDIs is not duplicated. This leads to a situation where trees of chained VDIs are created over time as VMs and their associated VDIs get cloned. When one of the VDIs in a chain is deleted, XenServer rationalizes the other VDIs in the chain to remove unnecessary VDIs.

This coalescing process runs asynchronously. The amount of disk space reclaimed and the time taken to perform the process depends on the size of the VDI and the amount of shared data. Only one coalescing process will ever be active for an SR. This process thread runs on the SR master host.

If you have critical VMs running on the master server of the pool and experience occasional slow IO due to this process, you can take steps to mitigate against this:

- Migrate the VM to a host other than the SR master
- Set the disk IO priority to a higher level, and adjust the scheduler. See [the section called "Virtual disk QoS settings"](#) for more information.

Space Utilization

Space utilization is always reported based on the current allocation of the SR, and may not reflect the amount of virtual disk space allocated. The reporting of space for LVM-based SRs versus File-based SRs will also differ given that File-based VHD supports full thin provisioning, while the underlying volume of an LVM-based VHD will be fully inflated to support potential growth for writeable leaf nodes. Space utilization



reported for the SR will depend on the number of snapshots, and the amount of difference data written to a disk between each snapshot.

LVM-based space utilization differs depending on whether an LVM SR is upgraded or created as a new SR in XenServer. Upgraded LVM SRs will retain a base node that is fully inflated to the size of the virtual disk, and any subsequent snapshot or clone operations will provision at least one additional node that is fully inflated. For new SRs, in contrast, the base node will be deflated to only the data allocated in the VHD overlay.

When VHD-based VDIs are deleted, the space is marked for deletion on disk. Actual removal of allocated data may take some time to occur as it is handled by the coalesce process that runs asynchronously and independently for each VHD-based SR.

LUN-based VDIs

Mapping a raw LUN as a Virtual Disk image is typically the most high-performance storage method. For administrators that want to leverage existing storage SAN infrastructure such as Netapp, Equallogic or StorageLink accessible arrays, the array snapshot, clone and thin provisioning capabilities can be exploited directly using one of the array specific adapter SR types (Netapp, Equallogic or StorageLink). The virtual machine storage operations are mapped directly onto the array APIs using a LUN per VDI representation. This includes activating the data path on demand such as when a VM is started or migrated to another host.

Managed NetApp LUNs are accessible using the NetApp SR driver type, and are hosted on a Network Appliance device running a version of Ontap 7.0 or greater. LUNs are allocated and mapped dynamically to the host using the XenServer host management framework.

EqualLogic storage is accessible using the EqualLogic SR driver type, and is hosted on an EqualLogic storage array running a firmware version of 4.0 or greater. LUNs are allocated and mapped dynamically to the host using the XenServer host management framework.

For further information on StorageLink supported array systems and the various capabilities in each case, please refer to the StorageLink documentation directly.

Storage configuration

This section covers creating storage repository types and making them available to a XenServer host. The examples provided pertain to storage configuration using the CLI, which provides the greatest flexibility. See the XenCenter Help for details on using the **New Storage Repository** wizard.

Creating Storage Repositories

This section explains how to create Storage Repositories (SRs) of different types and make them available to a XenServer host. The examples provided cover creating SRs using the xe CLI. See the XenCenter help for details on using the **New Storage Repository** wizard to add SRs using XenCenter.

Note:

Local SRs of type `lvm` and `ext` can only be created using the xe CLI. After creation all SR types can be managed by either XenCenter or the xe CLI.

There are two basic steps involved in creating a new storage repository for use on a XenServer host using the CLI:

1. Probe the SR type to determine values for any required parameters.
2. Create the SR to initialize the SR object and associated PBD objects, plug the PBDs, and activate the SR.

These steps differ in detail depending on the type of SR being created. In all examples the **sr-create** command returns the UUID of the created SR if successful.



SRs can also be *destroyed* when no longer in use to free up the physical device, or *forgotten* to detach the SR from one XenServer host and attach it to another. See [the section called “Destroying or forgetting a SR”](#) for details.

Note:

When specifying StorageLink configuration for a XenServer host or pool, supply either the default credentials of username: **admin** and password: **storagelink**, or any custom credentials specified during installation of the StorageLink Gateway service. Unlike StorageLink Manager, XenCenter does not supply the default credentials automatically.

Upgrading LVM storage from XenServer 5.0 or earlier

See the *XenServer Installation Guide* for information on upgrading LVM storage to enable the latest features. Local, LVM on iSCSI, and LVM on HBA storage types from older (XenServer 5.0 and before) product versions will need to be upgraded before they will support snapshot and fast clone.

Warning:

SR upgrade of SRs created in version 5.0 or before requires the creation of a 4MB metadata volume. Please ensure that there are at least 4MB of free space on your SR before attempting to upgrade the storage.

Note:

Upgrade is a one-way operation so Citrix recommends only performing the upgrade when you are certain the storage will no longer need to be attached to a pool running an older software version.

LVM performance considerations

The snapshot and fast clone functionality provided in XenServer 5.5 and later for LVM-based SRs comes with an inherent performance overhead. In cases where optimal performance is desired, XenServer supports creation of VDIs in the *raw* format in addition to the default VHD format. The XenServer snapshot functionality is not supported on raw VDIs.

Note:

Non-transportable snapshots using the default Windows VSS provider will work on any type of VDI.

Warning:

Do not try to snapshot a VM that has `type=raw` disks attached. This could result in a partial snapshot being created. In this situation, you can identify the orphan snapshot VDIs by checking the `snapshot-of` field and then deleting them.

VDI types

In general, VHD format VDIs will be created. You can opt to use raw at the time you create the VDI; this can only be done using the `xe` CLI. After software upgrade from a previous XenServer version, existing data will be preserved as backwards-compatible raw VDIs but these are special-cased so that snapshots can be taken of them once you have allowed this by upgrading the SR. Once the SR has been upgraded and the first snapshot has been taken, you will be accessing the data through a VHD format VDI.

To check if an SR has been upgraded, verify that its `sm-config:use_vhd` key is `true`. To check if a VDI was created with `type=raw`, check its `sm-config` map. The **sr-param-list** and **vgi-param-list** `xe` commands can be used respectively for this purpose.

Creating a raw virtual disk using the `xe` CLI

1. Run the following command to create a VDI given the UUID of the SR you want to place the virtual disk in:



```
xe vdi-create sr-uuid=<sr-uuid> type=user virtual-size=<virtual-size> \
  name-label=<VDI name> sm-config:type=raw
```

2. Attach the new virtual disk to a VM and use your normal disk tools within the VM to partition and format, or otherwise make use of the new disk. You can use the **vbd-create** command to create a new VBD to map the virtual disk into your VM.

Converting between VDI formats

It is not possible to do a direct conversion between the raw and VHD formats. Instead, you can create a new VDI (either raw, as described above, or VHD if the SR has been upgraded or was created on XenServer 5.5 or later) and then copy data into it from an existing volume. Citrix recommends that you use the xe CLI to ensure that the new VDI has a virtual size at least as big as the VDI you are copying from (by checking its virtual-size field, for example by using the **vdi-param-list** command). You can then attach this new VDI to a VM and use your preferred tool within the VM (standard disk management tools in Windows, or the **dd** command in Linux) to do a direct block-copy of the data. If the new volume is a VHD volume, it is important to use a tool that can avoid writing empty sectors to the disk so that space is used optimally in the underlying storage repository — in this case a file-based copy approach may be more suitable.

Probing an SR

The **sr-probe** command can be used in two ways:

1. To identify unknown parameters for use in creating a SR.
2. To return a list of existing SRs.

In both cases **sr-probe** works by specifying an SR type and one or more `device-config` parameters for that SR type. When an incomplete set of parameters is supplied the **sr-probe** command returns an error message indicating parameters are missing and the possible options for the missing parameters. When a complete set of parameters is supplied a list of existing SRs is returned. All **sr-probe** output is returned as XML.

For example, a known iSCSI target can be probed by specifying its name or IP address, and the set of IQNs available on the target will be returned:

```
xe sr-probe type=lvmoiscsi device-config:target=<192.168.1.10>
```

Error code: SR_BACKEND_FAILURE_96

Error parameters: , The request is missing or has an incorrect target IQN parameter, \

```
<?xml version="1.0" ?>
```

```
<iscsi-target-iqns>
  <TGT>
    <Index>
      0
    </Index>
    <IPAddress>
      192.168.1.10
    </IPAddress>
    <TargetIQN>
      iqn.192.168.1.10:filer1
    </TargetIQN>
  </TGT>
</iscsi-target-iqns>
```

Probing the same target again and specifying both the name/IP address and desired IQN returns the set of SCSI LUNs (LUNs) available on the target/IQN.



```
xe sr-probe type=lvmoiscsi device-config:target=192.168.1.10 \
device-config:targetIQN=iqn.192.168.1.10:filer1
```

Error code: SR_BACKEND_FAILURE_107

Error parameters: , The SCSIid parameter is missing or incorrect, \

```
<?xml version="1.0" ?>
```

```
<iscsi-target>
```

```
  <LUN>
```

```
    <vendor>
```

```
    IET
```

```
    </vendor>
```

```
    <LUNid>
```

```
      0
```

```
    </LUNid>
```

```
    <size>
```

```
      42949672960
```

```
    </size>
```

```
    <SCSIid>
```

```
      149455400000000000000000002000000b70200000f000000
```

```
    </SCSIid>
```

```
  </LUN>
```

```
</iscsi-target>
```

Probing the same target and supplying all three parameters will return a list of SRs that exist on the LUN, if any.

```
xe sr-probe type=lvmoiscsi device-config:target=192.168.1.10 \
device-config:targetIQN=192.168.1.10:filer1 \
device-config:SCSIid=149455400000000000000000002000000b70200000f000000
```

```
<?xml version="1.0" ?>
```

```
<SRlist>
```

```
  <SR>
```

```
    <UUID>
```

```
      3f6e1ebd-8687-0315-f9d3-b02ab3adc4a6
```

```
    </UUID>
```

```
    <Devlist>
```

```
      /dev/disk/by-id/scsi-149455400000000000000000002000000b70200000f000000
```

```
    </Devlist>
```

```
  </SR>
```

```
</SRlist>
```

The following parameters can be probed for each SR type:

SR type	device-config parameter, in order of dependency	Can be probed?	Required for sr-create?
lvmoiscsi	target	No	Yes
	chapuser	No	No
	chappassword	No	No
	targetIQN	Yes	Yes
	SCSIid	Yes	Yes
lvmohba	SCSIid	Yes	Yes

SR type	device-config parameter, in order of dependency	Can be probed?	Required for sr-create?
netapp	target	No	Yes
	username	No	Yes
	password	No	Yes
	chapuser	No	No
	chappassword	No	No
	aggregate	No [*]	Yes
	FlexVols	No	No
	allocation	No	No
	asis	No	No
nfs	server	No	Yes
	serverpath	Yes	Yes
lvm	device	No	Yes
ext	device	No	Yes
equallogic	target	No	Yes
	username	No	Yes
	password	No	Yes
	chapuser	No	No
	chappassword	No	No
	storagepool	No [†]	Yes
cslg	target	No	Yes
	storageSystemId	Yes	Yes
	storagePoolId	Yes	Yes
	username	No	No [‡]
	password	No	No [‡]
	cslport	No	No [‡]
	chapuser	No	No [‡]
	chappassword	No	No [‡]

SR type	device-config parameter, in order of dependency	Can be probed?	Required for sr-create?
	provision-type	Yes	No
	protocol	Yes	No
	provision-options	Yes	No
	raid-type	Yes	No

*Aggregate probing is only possible at **sr-create** time. It needs to be done there so that the aggregate can be specified at the point that the SR is created.

†Storage pool probing is only possible at **sr-create** time. It needs to be done there so that the aggregate can be specified at the point that the SR is created.

‡If the username, password, or port configuration of the StorageLink service are changed from the default value then the appropriate parameter and value must be specified.

Storage Multipathing

Dynamic multipathing support is available for Fibre Channel and iSCSI storage backends. By default, it uses round-robin mode load balancing, so both routes have active traffic on them during normal operation. You can enable multipathing in XenCenter or on the xe CLI.

Before attempting to enable multipathing, verify that multiple targets are available on your storage server. For example, an iSCSI storage backend queried for `sendtargets` on a given portal should return multiple targets, as in the following example:

```
iscsiadm -m discovery --type sendtargets --portal 192.168.0.161
192.168.0.161:3260,1 iqn.strawberry:litchie
192.168.0.204:3260,2 iqn.strawberry:litchie
```

To enable storage multipathing using the xe CLI

1. Unplug all PBDs on the host:

```
xe pbd-unplug uuid=<pbid_uuid>
```

2. Set the host's `other-config:multipathing` parameter:

```
xe host-param-set other-config:multipathing=true uuid=host_uuid
```

3. Set the host's `other-config:multipathhandle` parameter to `dmp`:

```
xe host-param-set other-config:multipathhandle=dmp uuid=host_uuid
```

4. If there are existing SRs on the host running in single path mode but that have multiple paths:

- Migrate or suspend any running guests with virtual disks in affected the SRs
- Unplug and re-plug the PBD of any affected SRs to reconnect them using multipathing:

```
xe pbd-plug uuid=<pbid_uuid>
```

To disable multipathing, first unplug your VBDs, set the host `other-config:multipathing` parameter to `false` and then replug your PBDs as described above. Do not modify the `other-config:multipathhandle` parameter as this will be done automatically.

Multipath support in XenServer is based on the device-mapper `multipathd` components. Activation and deactivation of multipath nodes is handled automatically by the Storage Manager API. Unlike the standard

`dm-multipath` tools in Linux, device mapper nodes are not automatically created for all LUNs on the system, and it is only when LUNs are actively used by the storage management layer that new device mapper nodes are provisioned. It is unnecessary therefore to use any of the `dm-multipath` CLI tools to query or refresh DM table nodes in XenServer. Should it be necessary to query the status of device-mapper tables manually, or list active device mapper multipath nodes on the system, use the `mpathutil` utility:

- `mpathutil list`
- `mpathutil status`

Unlike the standard `dm-multipath` tools in Linux, device mapper nodes are not automatically created for all LUNs on the system. As LUNs are actively used by the storage management layer, new device mapper nodes are provisioned. It is unnecessary to use any of the `dm-multipath` CLI tools to query or refresh DM table nodes in XenServer.

Note:

Due to incompatibilities with the integrated multipath management architecture, the standard `dm-multipath` CLI utility *should not be used* with XenServer. Please use the `mpathutil` CLI tool for querying the status of nodes on the host.

Note:

Multipath support in Equallogic arrays does not encompass Storage IO multipathing in the traditional sense of the term. Multipathing must be handled at the network/NIC bond level. Refer to the Equallogic documentation for information about configuring network failover for Equallogic SRs/LVMoISCSI SRs.

Storage Repository Types

The storage repository types supported in XenServer are provided by plugins in the control domain; these can be examined and plugins supported third parties can be added to the `/opt/xensource/sm` directory. Modification of these files is unsupported, but visibility of these files may be valuable to developers and power users. New storage manager plugins placed in this directory are automatically detected by XenServer. Use the `sm-list` command (see [the section called “Storage Manager commands”](#)) to list the available SR types.

New storage repositories are created using the **New Storage wizard** in XenCenter. The wizard guides you through the various probing and configuration steps. Alternatively, use the `sr-create` command. This command creates a new SR on the storage substrate (potentially destroying any existing data), and creates the SR API object and a corresponding PBD record, enabling VMs to use the storage. On successful creation of the SR, the PBD is automatically plugged. If the SR `shared=true` flag is set, a PBD record is created and plugged for every XenServer Host in the resource pool.

All XenServer SR types support VDI resize, fast cloning and snapshot. SRs based on the LVM SR type (local, iSCSI, or HBA) provide thin provisioning for snapshot and hidden parent nodes. The other SR types support full thin provisioning, including for virtual disks that are active.

Note:

Automatic LVM metadata archiving is disabled by default. This does not prevent metadata recovery for LVM groups.

Warning:

When VHD VDIs are not attached, for example in the case of a VDI snapshot, they are stored by default thinly-provisioned. Because of this it is imperative to ensure that there is sufficient disk-space available for the VDI to become thickly provisioned when attempting to attach it. VDI clones, however, are thickly-provisioned.



The maximum supported VDI sizes are:

Storage type	Maximum VDI size
EXT3	2TB
LVM	2TB
Netapp	2TB
EqualLogic	15TB
ONTAP(NetApp)	12TB

Local LVM

The Local LVM type presents disks within a locally-attached Volume Group.

By default, XenServer uses the local disk on the physical host on which it is installed. The Linux Logical Volume Manager (LVM) is used to manage VM storage. A VDI is implemented in VHD format in an LVM logical volume of the specified size.

XenServer versions prior to 5.6 did not use the VHD format and will remain in legacy mode. See [the section called “Upgrading LVM storage from XenServer 5.0 or earlier”](#) for information about upgrading a storage repository to the new format.

Creating a local LVM SR (lvm)

Device-config parameters for lvm SRs are:

Parameter Name	Description	Required?
Device	device name on the local host to use for the SR	Yes

To create a local lvm SR on /dev/sdb use the following command.

```
xe sr-create host-uuid=<valid_uuid> content-type=user \  
name-label=<"Example Local LVM SR"> shared=false \  
device-config:device=/dev/sdb type=lvm
```

Local EXT3 VHD

The Local EXT3 VHD type represents disks as VHD files stored on a local path.

Local disks can also be configured with a local EXT SR to serve VDIs stored in the VHD format. Local disk EXT SRs must be configured using the XenServer CLI.

By definition, local disks are not shared across pools of XenServer host. As a consequence, VMs whose VDIs are stored in SRs on local disks are not agile -- they cannot be migrated between XenServer hosts in a resource pool.

Creating a local EXT3 SR (ext)

Device-config parameters for ext SRs:

Parameter Name	Description	Required?
Device	device name on the local host to use for the SR	Yes

To create a local ext SR on `/dev/sdb` use the following command:

```
xe sr-create host-uuid=<valid_uuid> content-type=user \
name-label=<"Example Local EXT3 SR"> shared=false \
device-config:device=/dev/sdb type=ext
```

udev

The `udev` type represents devices plugged in using the `udev` device manager as VDIs.

XenServer has two SRs of type `udev` that represent removable storage. One is for the CD or DVD disk in the physical CD or DVD-ROM drive of the XenServer host. The other is for a USB device plugged into a USB port of the XenServer host. VDIs that represent the media come and go as disks or USB sticks are inserted and removed.

ISO

The ISO type handles CD images stored as files in ISO format. This SR type is useful for creating shared ISO libraries. For storage repositories that store a library of ISOs, the `content-type` parameter must be set to `iso`.

For example:

```
xe sr-create host-uuid=<valid_uuid> content-type=iso \
type=iso name-label=<"Example ISO SR"> location=<nfs server:path>>
```

EqualLogic

The EqualLogic SR type maps LUNs to VDIs on an EqualLogic array group, allowing for the use of fast snapshot and clone features on the array.

If you have access to an EqualLogic filer, you can configure a custom EqualLogic storage repository for VM storage on your XenServer deployment. This allows the use of the advanced features of this filer type. Virtual disks are stored on the filer using one LUN per virtual disk. Using this storage type will enable the thin provisioning, snapshot, and fast clone features of this filer.

Consider your storage requirements when deciding whether to use the specialized SR plugin, or to use the generic LVM/iSCSI storage backend. By using the specialized plugin, XenServer will communicate with the filer to provision storage. Some arrays have a limitation of seven concurrent connections, which may limit the throughput of control operations. Using the plugin will allow you to make use of the advanced array features, however, so will make backup and snapshot operations easier.

Warning:

There are two types of administration accounts that can successfully access the EqualLogic SM plugin:

- A group administration account which has access to and can manage the entire group and all storage pools.
- A pool administrator account that can manage only the objects (SR and VDI snapshots) that are in the pool or pools assigned to the account.

Creating a shared EqualLogic SR

Device-config parameters for EqualLogic SRs:

Parameter Name	Description	Optional?
target	the IP address or hostname of the EqualLogic array that hosts the SR	no
username	the login username used to manage the LUNs on the array	no
password	the login password used to manage the LUNs on the array	no
storagepool	the storage pool name	no
chapuser	the username to be used for CHAP authentication	yes
chappassword	the password to be used for CHAP authentication	yes
allocation	specifies whether to use thick or thin provisioning. Default is thick. Thin provisioning reserves a minimum of 10% of volume space.	yes
snap-reserve-percentage	sets the amount of space, as percentage of volume reserve, to allocate to snapshots. Default is 100%.	yes
snap-depletion	sets the action to take when snapshot reserve space is exceeded. <i>volume-offline</i> sets the volume and all its snapshots offline. This is the default action. The <i>delete-oldest</i> action deletes the oldest snapshot until enough space is available for creating the new snapshot.	yes
control	Certain customer configurations may require separate IP addresses for the control and iSCSI target interfaces; use this option to specify a different control IP address from the device-config target address.	yes

Use the **sr-create** command to create an EqualLogic SR. For example:

```
xe sr-create host-uuid=<valid_uuid> content-type=user \
name-label=<"Example shared Equallogic SR"> \
shared=true device-config:target=<target_ip> \
device-config:username=<admin_username> \
device-config:password=<admin_password> \
device-config:storagepool=<my_storagepool> \
device-config:chapuser=<chapusername> \
device-config:chappassword=<chapuserpassword> \
device-config:allocation=<thick> \
type=equal
```

EqualLogic VDI Snapshot space allocation with XenServer EqualLogic Adapter

When you create a SR using the EqualLogic plug-in, you specify a storage pool in which the SR is created. This assumes that the free space in the storage pool will be used for creating the VDIs, and for snapshot and clones when requested. If the storage pool comprises of all the "member arrays" in the EqualLogic group



then the plug-in will use all of the space on the SAN for creating VDIs. When the SR is created, a small amount of meta data is created, called the SR Management Volume. This will be displayed as the smallest volume (30MB). All of the VDIs in the SR are created with 100% space reserved for snapshots by default. This will also override the Group Defaults that are set by the administrator from the EqualLogic GUI/CLI. You can change this default by using a `sm-config` parameter while creating a VDI via the xe CLI.

Creating a VDI using the CLI

To create a VDI using CLI use the `xe vdi-create` command:

```
xe vdi-create virtual-size=10GiB name-label=testvdi \  
    sr-uuid=<sr_uuid> type=user sm-config:allocation=<thin | thick> \  
    sm-config:snap-reserve-percentage=<Percentage Integral Value> \  
    sm-config:snap-depletion=<delete-oldest | volume-offline>
```

Where `<sr_uuid>` is the UUID of the SR of type Dell EqualLogic.

`sm-config:allocation` controls whether the VDI volume is provisioned as a thin volume or not. Setting `sm-config:allocation=thin` will create a volume with thin provisioning enabled. Setting `sm-config:allocation=thick` will create a volume with thick provisioning disabled. If the type of allocation is not specified, the default allocation for the SR is used to provision the VDI volume.

`sm-config:snap-reserve-percentage` specifies the amount of space, in terms of percentage of volume, to reserve for volume snapshots.

`sm-config:snap-depletion` specifies the snapshot space recovery policy action taken when the space reserved for snapshots has been exceeded. Setting `sm-config:snap-depletion=delete-oldest` deletes the oldest snapshots until sufficient space is recovered (the default).

Setting `sm-config:snap-depletion=volume-offline` sets the volume and snapshots offline. Active iSCSI connections will be terminated before a snapshot is automatically deleted.

NetApp

The NetApp type maps LUNs to VDIs on a NetApp server, enabling the use of fast snapshot and clone features on the filer.

Note:

NetApp and EqualLogic SRs require XenServer Advanced edition or above to use the special integration with the NetApp and Dell EqualLogic SR types, but you can use them as ordinary iSCSI, FC, or NFS storage with free XenServer, without the benefits of direct control of hardware features. To find out about XenServer editions and how to upgrade, visit the Citrix website [here](#).

If you have access to Network Appliance™ (NetApp) storage with sufficient disk space, running a version of Data ONTAP 7G (version 7.0 or greater), you can configure a custom NetApp storage repository for VM storage on your XenServer deployment. The XenServer driver uses the ZAPI interface to the storage to create a group of *FlexVols* that correspond to an SR. VDIs are created as virtual LUNs on the storage, and attached to XenServer hosts using an iSCSI data path. There is a direct mapping between a VDI and a raw LUN that does not require any additional volume metadata. The NetApp SR is a managed volume and the VDIs are the LUNs within the volume. VM cloning uses the snapshotting and cloning capabilities of the storage for data efficiency and performance and to ensure compatibility with existing ONTAP management tools.

As with the iSCSI-based SR type, the NetApp driver also uses the built-in software initiator and its assigned host IQN, which can be modified by changing the value shown on the **General** tab when the storage repository is selected in XenCenter.



The easiest way to create NetApp SRs is to use XenCenter. See the XenCenter help for details. See [the section called “Creating a shared NetApp SR over iSCSI”](#) for an example of how to create them using the xe CLI.

FlexVols

NetApp uses FlexVols as the basic unit of manageable data. There are limitations that constrain the design of NetApp-based SRs. These are:

- maximum number of FlexVols per filer
- maximum number of LUNs per network port
- maximum number of snapshots per FlexVol

Precise system limits vary per filer type, however as a general guide, a FlexVol may contain up to 200 LUNs, and provides up to 255 snapshots. Because there is a one-to-one mapping of LUNs to VDIs, and because often a VM will have more than one VDI, the resource limitations of a single FlexVol can easily be reached. Also, the act of taking a snapshot includes snapshotting all the LUNs within a FlexVol and the VM clone operation indirectly relies on snapshots in the background as well as the VDI snapshot operation for backup purposes.

There are two constraints to consider when mapping the virtual storage objects of the XenServer host to the physical storage. To maintain space efficiency it makes sense to limit the number of LUNs per FlexVol, yet at the other extreme, to avoid resource limitations a single LUN per FlexVol provides the most flexibility. However, because there is a vendor-imposed limit of 200 or 500 FlexVols, per filer (depending on the NetApp model), this creates a limit of 200 or 500 VDIs per filer and it is therefore important to select a suitable number of FlexVols taking these parameters into account.

Given these resource constraints, the mapping of virtual storage objects to the Ontap storage system has been designed in the following manner. LUNs are distributed evenly across FlexVols, with the expectation of using VM UUIDs to opportunistically group LUNs attached to the same VM into the same FlexVol. This is a reasonable usage model that allows a snapshot of all the VDIs in a VM at one time, maximizing the efficiency of the snapshot operation.

An optional parameter you can set is the number of FlexVols assigned to the SR. You can use between 1 and 32 FlexVols; the default is 8. The trade-off in the number of FlexVols to the SR is that, for a greater number of FlexVols, the snapshot and clone operations become more efficient, because there are fewer VMs backed off the same FlexVol. The disadvantage is that more FlexVol resources are used for a single SR, where there is a typical system-wide limitation of 200 for some smaller filers.

Aggregates

When creating a NetApp driver-based SR, you select an appropriate *aggregate*. The driver can be probed for non-traditional type aggregates, that is, newer-style aggregates that support FlexVols, and lists all aggregates available and the unused disk space on each.

Note:

Aggregate probing is only possible at **sr-create** time so that the aggregate can be specified at the point that the SR is created, but is not probed by the **sr-probe** command.

Citrix strongly recommends that you configure an aggregate exclusively for use by XenServer storage, because space guarantees and allocation cannot be correctly managed if other applications are sharing the resource.

Thick or thin provisioning

When creating NetApp storage, you can also choose the type of space management used. By default, allocated space is thickly provisioned to ensure that VMs never run out of disk space and that all virtual allocation guarantees are fully enforced on the filer. Selecting thick provisioning ensures that whenever a VDI (LUN) is allocated on the filer, sufficient space is reserved to guarantee that it will never run out of space and consequently experience failed writes to disk. Due to the nature of the Ontap FlexVol space provisioning algorithms the best practice guidelines for the filer require that at least twice the LUN space is reserved to account for background snapshot data collection and to ensure that writes to disk are never blocked. In addition to the double disk space guarantee, Ontap also requires some additional space reservation for management of unique blocks across snapshots. The guideline on this amount is 20% above the reserved space. The space guarantees afforded by thick provisioning will reserve up to 2.4 times the requested virtual disk space.

The alternative allocation strategy is *thin provisioning*, which allows the administrator to present more storage space to the VMs connecting to the SR than is actually available on the SR. There are no space guarantees, and allocation of a LUN does not claim any data blocks in the FlexVol until the VM writes data. This might be appropriate for development and test environments where you might find it convenient to over-provision virtual disk space on the SR in the anticipation that VMs might be created and destroyed frequently without ever utilizing the full virtual allocated disk.

Warning:

If you are using thin provisioning in production environments, take appropriate measures to ensure that you never run out of storage space. VMs attached to storage that is full will fail to write to disk, and in some cases may fail to read from disk, possibly rendering the VM unusable.

FAS Deduplication

FAS Deduplication is a NetApp technology for reclaiming redundant disk space. Newly-stored data objects are divided into small blocks, each block containing a digital signature, which is compared to all other signatures in the data volume. If an exact block match exists, the duplicate block is discarded and the disk space reclaimed. FAS Deduplication can be enabled on thin provisioned NetApp-based SRs and operates according to the default filer FAS Deduplication parameters, typically every 24 hours. It must be enabled at the point the SR is created and any custom FAS Deduplication configuration must be managed directly on the filer.

Access Control

Because FlexVol operations such as volume creation and volume snapshotting require administrator privileges on the filer itself, Citrix recommends that the XenServer host is provided with suitable administrator username and password credentials at configuration time. In situations where the XenServer host does not have full administrator rights to the filer, the filer administrator could perform an out-of-band preparation and provisioning of the filer and then introduce the SR to the XenServer host using XenCenter or the **sr-introduce** xe CLI command. Note, however, that operations such as VM cloning or snapshot generation will fail in this situation due to insufficient access privileges.

Licenses

You need to have an iSCSI license on the NetApp filer to use this storage repository type; for the generic plugins you need either an iSCSI or NFS license depending on the SR type being used.

Further information

For more information about NetApp technology, see the following links:



- [General information on NetApp products](#)
- [Data ONTAP](#)
- [FlexVol](#)
- [FlexClone](#)
- [RAID-DP](#)
- [Snapshot](#)
- [FilerView](#)

Creating a shared NetApp SR over iSCSI

Device-config parameters for netapp SRs:

Parameter Name	Description	Optional?
target	the IP address or hostname of the NetApp server that hosts the SR	no
port	the port to use for connecting to the NetApp server that hosts the SR. Default is port 80.	yes
usehttps	specifies whether to use a secure TLS-based connection to the NetApp server that hosts the SR [true false]. Default is false.	yes
username	the login username used to manage the LUNs on the filer	no
password	the login password used to manage the LUNs on the filer	no
aggregate	the aggregate name on which the FlexVol is created	Required for sr_create
FlexVols	the number of FlexVols to allocate to each SR	yes
chapuser	the username for CHAP authentication	yes
chappassword	the password for CHAP authentication	yes
allocation	specifies whether to provision LUNs using thick or thin provisioning. Default is thick	yes
asis	specifies whether to use FAS Deduplication if available. Default is false	yes

Setting the SR *other-config:multiplier* parameter to a valid value adjusts the default multiplier attribute. By default XenServer allocates 2.4 times the requested space to account for snapshot and metadata overhead associated with each LUN. To save disk space, you can set the multiplier to a value ≥ 1 . Setting the multiplier should only be done with extreme care by system administrators who understand the space allocation constraints of the NetApp filer. If you try to set the amount to less than 1, for example, in an attempt to pre-allocate very little space for the LUN, the attempt will most likely fail.

Setting the SR *other-config:enforce_allocation* parameter to *true* resizes the FlexVols to precisely the amount specified by either the *multiplier* value above, or the default 2.4 value.



Note:

This works on new VDI creation in the selected FlexVol, or on all FlexVols during an SR scan and overrides any manual size adjustments made by the administrator to the SR FlexVols.

To create a NetApp SR, use the following command.

```
xe sr-create host-uuid=<valid_uuid> content-type=user \  
  name-label=<"Example shared NetApp SR"> shared=true \  
  device-config:target=<192.168.1.10> device-config:username=<admin_username> \  
  device-config:password=<admin_password> \  
  type=netapp
```

Managing VDIs in a NetApp SR

Due to the complex nature of mapping VM storage objects onto NetApp storage objects such as LUNs, FlexVols and disk Aggregates, the plugin driver makes some general assumptions about how storage objects should be organized. The default number of FlexVols that are managed by an SR instance is 8, named `XenStorage_<SR_UUID>_FV<#>` where # is a value between 0 and the total number of FlexVols assigned. This means that VDIs (LUNs) are evenly distributed across any one of the FlexVols at the point that the VDI is instantiated. The only exception to this rule is for groups of VM disks which are opportunistically assigned to the same FlexVol to assist with VM cloning, and when VDIs are created manually but passed a `vmhint` flag that informs the backend of the FlexVol to which the VDI should be assigned. The `vmhint` may be a random string, such as a uuid that is re-issued for all subsequent VDI creation operations (to ensure grouping in the same FlexVol), or it can be a simple FlexVol number to correspond to the FlexVol naming convention applied on the Filer. Using either of the following 2 commands, a VDI created manually using the CLI can be assigned to a specific FlexVol:

```
xe vdi-create uuid=<valid_vdi_uuid> sr-uuid=<valid_sr_uuid> \  
  sm-config:vmhint=<valid_vm_uuid>  
  
xe vdi-create uuid=<valid_vdi_uuid> sr-uuid=<valid_sr_uuid> \  
  sm-config:vmhint=<valid_flexvol_number>
```

Taking VDI snapshots with a NetApp SR

Cloning a VDI entails generating a snapshot of the FlexVol and then creating a LUN clone backed off the snapshot. When generating a VM snapshot you must snapshot each of the VMs disks in sequence. Because all the disks are expected to be located in the same FlexVol, and the FlexVol snapshot operates on all LUNs in the same FlexVol, it makes sense to re-use an existing snapshot for all subsequent LUN clones. By default, if no snapshot hint is passed into the backend driver it will generate a random ID with which to name the FlexVol snapshot. There is a CLI override for this value, passed in as an `epochhint`. The first time the `epochhint` value is received, the backend generates a new snapshot based on the cookie name. Any subsequent snapshot requests with the same `epochhint` value will be backed off the existing snapshot:

```
xe vdi-snapshot uuid=<valid_vdi_uuid> driver-params:epochhint=<cookie>
```

During NetApp SR provisioning, additional disk space is reserved for snapshots. If you plan to not use the snapshotting functionality, you might want to free up this reserved space. To do so, you can reduce the value of the `other-config:multiplier` parameter. By default the value of the multiplier is 2.4, so the amount of space reserved is 2.4 times the amount of space that would be needed for the FlexVols themselves.

Software iSCSI Support

XenServer provides support for shared SRs on iSCSI LUNs. iSCSI is supported using the open-iSCSI software iSCSI initiator or by using a supported iSCSI Host Bus Adapter (HBA). The steps for using iSCSI



HBAs are identical to those for Fibre Channel HBAs, both of which are described in [the section called “Creating a shared LVM over Fibre Channel / iSCSI HBA or SAS SR \(lvmohba\)”](#).

Shared iSCSI support using the software iSCSI initiator is implemented based on the Linux Volume Manager (LVM) and provides the same performance benefits provided by LVM VDIs in the local disk case. Shared iSCSI SRs using the software-based host initiator are capable of supporting VM agility using XenMotion: VMs can be started on any XenServer host in a resource pool and migrated between them with no noticeable downtime.

iSCSI SRs use the entire LUN specified at creation time and may not span more than one LUN. CHAP support is provided for client authentication, during both the data path initialization and the LUN discovery phases.

XenServer Host iSCSI configuration

All iSCSI initiators and targets must have a unique name to ensure they can be uniquely identified on the network. An initiator has an iSCSI initiator address, and a target has an iSCSI target address. Collectively these are called iSCSI Qualified Names, or IQNs.

XenServer hosts support a single iSCSI initiator which is automatically created and configured with a random IQN during host installation. The single initiator can be used to connect to multiple iSCSI targets concurrently.

iSCSI targets commonly provide access control using iSCSI initiator IQN lists, so all iSCSI targets/LUNs to be accessed by a XenServer host must be configured to allow access by the host's initiator IQN. Similarly, targets/LUNs to be used as shared iSCSI SRs must be configured to allow access by all host IQNs in the resource pool.

Note:

iSCSI targets that do not provide access control will typically default to restricting LUN access to a single initiator to ensure data integrity. If an iSCSI LUN is intended for use as a shared SR across multiple XenServer hosts in a resource pool, ensure that multi-initiator access is enabled for the specified LUN.

The XenServer host IQN value can be adjusted using XenCenter, or using the CLI with the following command when using the iSCSI software initiator:

```
xe host-param-set uuid=<valid_host_id> other-config:iscsi_ign=<new_initiator_ign>
```

Warning:

It is imperative that every iSCSI target and initiator have a unique IQN. If a non-unique IQN identifier is used, data corruption and/or denial of LUN access can occur.

Warning:

Do not change the XenServer host IQN with iSCSI SRs attached. Doing so can result in failures connecting to new targets or existing SRs.

Managing Hardware Host Bus Adapters (HBAs)

This section covers various operations required to manage SAS, Fibre Channel and iSCSI HBAs.

Sample QLogic iSCSI HBA setup

For full details on configuring QLogic Fibre Channel and iSCSI HBAs please refer to the [QLogic website](#).



Once the HBA is physically installed into the XenServer host, use the following steps to configure the HBA:

1. Set the IP networking configuration for the HBA. This example assumes DHCP and HBA port 0. Specify the appropriate values if using static IP addressing or a multi-port HBA.

```
/opt/QLLogic_Corporation/SANsurferiCLI/iscli -ipdhcp 0
```

2. Add a persistent iSCSI target to port 0 of the HBA.

```
/opt/QLLogic_Corporation/SANsurferiCLI/iscli -pa 0 <iscsi_target_ip_address>
```

3. Use the **xe sr-probe** command to force a rescan of the HBA controller and display available LUNs. See [the section called "Probing an SR"](#) and [the section called "Creating a shared LVM over Fibre Channel / iSCSI HBA or SAS SR \(lvmohba\)"](#) for more details.

Removing HBA-based SAS, FC or iSCSI device entries

Note:

This step is not required. Citrix recommends that only power users perform this process if it is necessary.

Each HBA-based LUN has a corresponding global device path entry under `/dev/disk/by-scsibus` in the format `<SCSIid>-<adapter>:<bus>:<target>:<lun>` and a standard device path under `/dev`. To remove the device entries for LUNs no longer in use as SRs use the following steps:

1. Use **sr-forget** or **sr-destroy** as appropriate to remove the SR from the XenServer host database. See [the section called "Destroying or forgetting a SR"](#) for details.
2. Remove the zoning configuration within the SAN for the desired LUN to the desired host.
3. Use the **sr-probe** command to determine the ADAPTER, BUS, TARGET, and LUN values corresponding to the LUN to be removed. See [the section called "Probing an SR"](#) for details.
4. Remove the device entries with the following command:

```
echo "1" > /sys/class/scsi_device/<adapter>:<bus>:<target>:<lun>/device/delete
```

Warning:

Make absolutely sure you are certain which LUN you are removing. Accidentally removing a LUN required for host operation, such as the boot or root device, will render the host unusable.

LVM over iSCSI

The LVM over iSCSI type represents disks as Logical Volumes within a Volume Group created on an iSCSI LUN.

Creating a shared LVM over iSCSI SR using the software iSCSI initiator (lvmoiscsi)

Device-config parameters for lvmoiscsi SRs:

Parameter Name	Description	Optional
target	the IP address or hostname of the iSCSI filer that hosts the SR	yes
targetIQN	the IQN target address of iSCSI filer that hosts the SR	yes
SCSIid	the SCSI bus ID of the destination LUN	yes

Parameter Name	Description	Optional
chapuser	the username to be used for CHAP authentication	no
chappassword	the password to be used for CHAP authentication	no
port	the network port number on which to query the target	no
usediscoverynumber	the specific iscsi record index to use	no

To create a shared lvmoiscsi SR on a specific LUN of an iSCSI target use the following command.

```
xe sr-create host-uuid=<valid_uuid> content-type=user \
name-label=<"Example shared LVM over iSCSI SR"> shared=true \
device-config:target=<target_ip=> device-config:targetIQN=<target_iqn=> \
device-config:SCSIid=<scsci_id> \
type=lvmoiscsi
```

Creating a shared LVM over Fibre Channel / iSCSI HBA or SAS SR (lvmohba)

SRs of type lvmohba can be created and managed using the xe CLI or XenCenter.

Device-config parameters for lvmohba SRs:

Parameter name	Description	Required?
SCSIid	Device SCSI ID	Yes

To create a shared lvmohba SR, perform the following steps on each host in the pool:

1. Zone in one or more LUNs to each XenServer host in the pool. This process is highly specific to the SAN equipment in use. Please refer to your SAN documentation for details.
2. If necessary, use the HBA CLI included in the XenServer host to configure the HBA:
 - Emulex: /usr/sbin/hbanyware
 - QLogic FC: /opt/QLogic_Corporation/SANsurferCLI
 - QLogic iSCSI: /opt/QLogic_Corporation/SANsurferiCLI

See [the section called "Managing Hardware Host Bus Adapters \(HBAs\)"](#) for an example of QLogic iSCSI HBA configuration. For more information on Fibre Channel and iSCSI HBAs please refer to the [Emulex](#) and [QLogic](#) websites.

3. Use the **sr-probe** command to determine the global device path of the HBA LUN. **sr-probe** forces a re-scan of HBAs installed in the system to detect any new LUNs that have been zoned to the host and returns a list of properties for each LUN found. Specify the `host-uuid` parameter to ensure the probe occurs on the desired host.

The global device path returned as the `<path>` property will be common across all hosts in the pool and therefore must be used as the value for the `device-config:device` parameter when creating the SR.

If multiple LUNs are present use the vendor, LUN size, LUN serial number, or the SCSI ID as included in the `<path>` property to identify the desired LUN.



```
xe sr-probe type=lvmoaha \  
host-uuid=1212c7b3-f333-4a8d-a6fb-80c5b79b5b31  
Error code: SR_BACKEND_FAILURE_90  
Error parameters: , The request is missing the device parameter, \  
<?xml version="1.0" ?>  
<Devlist>  
  <BlockDevice>  
    <path>  
      /dev/disk/by-id/scsi-360a9800068666949673446387665336f  
    </path>  
    <vendor>  
      HITACHI  
    </vendor>  
    <serial>  
      730157980002  
    </serial>  
    <size>  
      80530636800  
    </size>  
    <adapter>  
      4  
    </adapter>  
    <channel>  
      0  
    </channel>  
    <id>  
      4  
    </id>  
    <lun>  
      2  
    </lun>  
    <hba>  
      qla2xxx  
    </hba>  
  </BlockDevice>  
  <Adapter>  
    <host>  
      Host4  
    </host>  
    <name>  
      qla2xxx  
    </name>  
    <manufacturer>  
      QLogic HBA Driver  
    </manufacturer>  
    <id>  
      4  
    </id>  
  </Adapter>  
</Devlist>
```

4. On the master host of the pool create the SR, specifying the global device path returned in the `<path>` property from **sr-probe**. PBDs will be created and plugged for each host in the pool automatically.

```
xe sr-create host-uuid=<valid_uuid> \  
content-type=user \  
name-label=<"Example shared LVM over HBA SR"> shared=true \  
device-config:SCSIid=<device_scsi_id> type=lvmoaha
```

Note:

You can use the `BRAND_CONSOLE`; **Repair Storage Repository** function to retry the PBD creation and plugging portions of the **sr-create** operation. This can be valuable in cases where the LUN zoning was incorrect for one or more hosts in a pool when the SR was created. Correct the zoning for the affected hosts and use the **Repair Storage Repository** function instead of removing and re-creating the SR.

NFS VHD

The NFS VHD type stores disks as VHD files on a remote NFS filesystem.

NFS is a ubiquitous form of storage infrastructure that is available in many environments. XenServer allows existing NFS servers that support NFS V3 over TCP/IP to be used immediately as a storage repository for virtual disks (VDIs). VDIs are stored in the Microsoft VHD format only. Moreover, as NFS SRs can be shared, VDIs stored in a shared SR allow VMs to be started on any XenServer hosts in a resource pool and be migrated between them using XenMotion with no noticeable downtime.

Creating an NFS SR requires the hostname or IP address of the NFS server. The **sr-probe** command provides a list of valid destination paths exported by the server on which the SR can be created. The NFS server must be configured to export the specified path to all XenServer hosts in the pool, or the creation of the SR and the plugging of the PBD record will fail.

As mentioned at the beginning of this chapter, VDIs stored on NFS are sparse. The image file is allocated as the VM writes data into the disk. This has the considerable benefit that VM image files take up only as much space on the NFS storage as is required. If a 100GB VDI is allocated for a new VM and an OS is installed, the VDI file will only reflect the size of the OS data that has been written to the disk rather than the entire 100GB.

VHD files may also be chained, allowing two VDIs to share common data. In cases where a NFS-based VM is cloned, the resulting VMs will share the common on-disk data at the time of cloning. Each will proceed to make its own changes in an isolated copy-on-write version of the VDI. This feature allows NFS-based VMs to be quickly cloned from templates, facilitating very fast provisioning and deployment of new VMs.

Note:

The maximum supported length of VHD chains is 30.

As VHD-based images require extra metadata to support sparseness and chaining, the format is not as high-performance as LVM-based storage. In cases where performance really matters, it is well worth forcibly allocating the sparse regions of an image file. This will improve performance at the cost of consuming additional disk space.

XenServer's NFS and VHD implementations assume that they have full control over the SR directory on the NFS server. Administrators should not modify the contents of the SR directory, as this can risk corrupting the contents of VDIs.

XenServer has been tuned for enterprise-class storage that use non-volatile RAM to provide fast acknowledgments of write requests while maintaining a high degree of data protection from failure. XenServer has been tested extensively against Network Appliance FAS270c and FAS3020c storage, using Data OnTap 7.2.2.

In situations where XenServer is used with lower-end storage, it will cautiously wait for all writes to be acknowledged before passing acknowledgments on to guest VMs. This will incur a noticeable performance cost, and might be remedied by setting the storage to present the SR mount point as an asynchronous



mode export. Asynchronous exports acknowledge writes that are not actually on disk, and so administrators should consider the risks of failure carefully in these situations.

The XenServer NFS implementation uses TCP by default. If your situation allows, you can configure the implementation to use UDP in situations where there may be a performance benefit. To do this, specify the *device-config* parameter `useUDP=true` at SR creation time.

Warning:

Since VDI on NFS SRs are created as sparse, administrators must ensure that there is enough disk space on the NFS SRs for all required VDIs. XenServer hosts do not enforce that the space required for VDIs on NFS SRs is actually present.

Creating a shared NFS SR (nfs)

Device-config parameters for nfs SRs:

Parameter Name	Description	Required?
server	IP address or hostname of the NFS server	Yes
serverpath	path, including the NFS mount point, to the NFS server that hosts the SR	Yes

To create a shared NFS SR on `192.168.1.10:/export1` use the following command.

```
xe sr-create host-uuid=<host_uuid> content-type=user \  
name-label=<"Example shared NFS SR"> shared=true \  
device-config:server=<192.168.1.10> device-config:serverpath=</export1> type=nfs
```

LVM over hardware HBA

The LVM over hardware HBA type represents disks as VHDs on Logical Volumes within a Volume Group created on an HBA LUN providing, for example, hardware-based iSCSI or FC support.

XenServer hosts support Fibre Channel (FC) storage area networks (SANs) through Emulex or QLogic host bus adapters (HBAs). All FC configuration required to expose a FC LUN to the host must be completed manually, including storage devices, network devices, and the HBA within the XenServer host. Once all FC configuration is complete the HBA will expose a SCSI device backed by the FC LUN to the host. The SCSI device can then be used to access the FC LUN as if it were a locally attached SCSI device.

Use the **sr-probe** command to list the LUN-backed SCSI devices present on the host. This command forces a scan for new LUN-backed SCSI devices. The path value returned by **sr-probe** for a LUN-backed SCSI device is consistent across all hosts with access to the LUN, and therefore must be used when creating shared SRs accessible by all hosts in a resource pool.

The same features apply to QLogic iSCSI HBAs.

See [the section called "Creating Storage Repositories"](#) for details on creating shared HBA-based FC and iSCSI SRs.

Note:



XenServer support for Fibre Channel does not support direct mapping of a LUN to a VM. HBA-based LUNs must be mapped to the host and specified for use in an SR. VDIs within the SR are exposed to VMs as standard block devices.

Citrix StorageLink Gateway (CSLG) SRs

The CSLG storage repository allows use of the Citrix StorageLink service for native access to a range of iSCSI and Fibre Channel arrays and automated fabric/initiator and array configuration features. Installation and configuration of the StorageLink service is required, for more information please see the StorageLink documentation.

Note:

Running the StorageLink service in a VM within a resource pool to which the StorageLink service is providing storage is not supported in combination with the XenServer High Availability (HA) features. To use CSLG SRs in combination with HA ensure the StorageLink service is running outside the HA-enabled pool.

CSLG SRs can be created using the xe CLI only. After creation CSLG SRs can be viewed and managed using both the xe CLI and XenCenter.

Because the CSLG SR can be used to access different storage arrays, the exact features available for a given CSLG SR depend on the capabilities of the array. All CSLG SRs use a LUN-per-VDI model where a new LUN is provisioned for each virtual disk. (VDI).

CSLG SRs can co-exist with other SR types on the same storage array hardware, and multiple CSLG SRs can be defined within the same resource pool.

The StorageLink service can be configured using the StorageLink Manager or from within the XenServer control domain using the StorageLink Command Line Interface (CLI). To run the StorageLink (CLI) use the following command, where *<hostname>* is the name or IP address of the machine running the StorageLink service:

```
/opt/Citrix/StorageLink/bin/csl \  
server=<hostname>[:<port>][,<username>,<password>]
```

For more information about the StorageLink CLI please see the StorageLink documentation or use the **/opt/Citrix/StorageLink/bin/csl help** command.

Creating a shared StorageLink SR

SRs of type CSLG can only be created by using the xe Command Line Interface (CLI). Once created CSLG SRs can be managed using either XenCenter or the xe CLI.

The *device-config* parameters for CSLG SRs are:

Parameter name	Description	Optional?
target	The server name or IP address of the machine running the StorageLink service	No
storageSystemId	The storage system ID to use for allocating storage	No

Parameter name	Description	Optional?
storagePoolId	The storage pool ID within the specified storage system to use for allocating storage	No
username	The username to use for connection to the StorageLink service	Yes *
password	The password to use for connecting to the StorageLink service	Yes *
csLport	The port to use for connecting to the StorageLink service	Yes *
chapuser	The username to use for CHAP authentication	Yes
chappassword	The password to use for CHAP authentication	Yes
protocol	Specifies the storage protocol to use (fc or iscsi) for multi-protocol storage systems. If not specified fc is used if available, otherwise iscsi.	Yes
provision-type	Specifies whether to use thick or thin provisioning (thick or thin); default is thick	Yes
provision-options	Additional provisioning options: Set to <code>dedup</code> to use the de-duplication features supported by the storage system	Yes
raid-type	The level of RAID to use for the SR, as supported by the storage array	Yes

* If the username, password, or port configuration of the StorageLink service are changed from the default then the appropriate parameter and value must be specified.

SRs of type `csLg` support two additional parameters that can be used with storage arrays that support LUN grouping features, such as NetApp flexvols.

The `sm-config` parameters for CSLG SRs are:

Parameter name	Description	Optional?
pool-count	Creates the specified number of groups on the array, in which LUNs provisioned within the SR will be created	Yes

Parameter name	Description	Optional?
physical-size	The total size of the SR in MB. Each pool will be created with a size equal to physical-size divided by pool-count.	Yes *

*Required when specifying the `sm-config:pool-count` parameter

Note:

When a new NetApp SR is created using StorageLink, by default a single FlexVol is created for the SR that contains all LUNs created for the SR. To change this behavior and specify the number of FlexVols to create and the size of each FlexVol, use the `sm-config:pool-size` and `sm-config:physical-size` parameters. The `sm-config:pool-size` parameter specifies the number of FlexVols. The `sm-config:physical-size` parameter specifies the total size of all FlexVols to be created, so that each FlexVol will be of size `sm-config:physical-size` divided by `sm-config:pool-size`.

To create a CSLG SR

1. Install the StorageLink service onto a Windows host or virtual machine
2. Configure the StorageLink service with the appropriate storage adapters and credentials
3. Use the **sr-probe** command with the `device-config:target` parameter to identify the available storage system IDs



```
xe sr-probe type=cslg device-config:target=192.168.128.10
```

```
<csl__storageSystemInfoList>
  <csl__storageSystemInfo>
    <friendlyName>5001-4380-013C-0240</friendlyName>
    <displayName>HP EVA (5001-4380-013C-0240)</displayName>
    <vendor>HP</vendor>
    <model>EVA</model>
    <serialNum>50014380013C0240</serialNum>
    <storageSystemId>HP__EVA__50014380013C0240</storageSystemId>
    <systemCapabilities>
      <capabilities>PROVISIONING</capabilities>
      <capabilities>MAPPING</capabilities>
      <capabilities>MULTIPLE_STORAGE_POOLS</capabilities>
      <capabilities>DIFF_SNAPSHOT</capabilities>
      <capabilities>CLONE</capabilities>
    </systemCapabilities>
    <protocolSupport>
      <capabilities>FC</capabilities>
    </protocolSupport>
    <csl__snapshotMethodInfoList>
      <csl__snapshotMethodInfo>
        <name>5001-4380-013C-0240</name>
        <displayName></displayName>
        <maxSnapshots>16</maxSnapshots>
        <supportedNodeTypes>
          <nodeType>STORAGE_VOLUME</nodeType>
        </supportedNodeTypes>
        <snapshotTypeList>
          </snapshotTypeList>
        <snapshotCapabilities>
          </snapshotCapabilities>
        </csl__snapshotMethodInfo>
        <csl__snapshotMethodInfo>
          <name>5001-4380-013C-0240</name>
          <displayName></displayName>
          <maxSnapshots>16</maxSnapshots>
          <supportedNodeTypes>
            <nodeType>STORAGE_VOLUME</nodeType>
          </supportedNodeTypes>
          <snapshotTypeList>
            <snapshotType>DIFF_SNAPSHOT</snapshotType>
          </snapshotTypeList>
          <snapshotCapabilities>
            </snapshotCapabilities>
          </csl__snapshotMethodInfo>
        <csl__snapshotMethodInfo>
          <name>5001-4380-013C-0240</name>
          <displayName></displayName>
          <maxSnapshots>16</maxSnapshots>
          <supportedNodeTypes>
            <nodeType>STORAGE_VOLUME</nodeType>
          </supportedNodeTypes>
          <snapshotTypeList>
            <snapshotType>CLONE</snapshotType>
          </snapshotTypeList>
          <snapshotCapabilities>
            </snapshotCapabilities>
          </csl__snapshotMethodInfo>
        </csl__snapshotMethodInfoList>
      </csl__storageSystemInfo>
    </csl__storageSystemInfoList>
```



You can use `grep` to filter the `sr-probe` output to just the storage pool IDs

```
xe sr-probe type=cslg device-config:target=192.168.128.10 | grep storageSystemId
<storageSystemId>EMC__CLARIION__APM00074902515</storageSystemId>
<storageSystemId>HP__EVA__50014380013C0240</storageSystemId>
<storageSystemId>NETAPP__LUN__0AD4F00A</storageSystemId>
```

4. Add the desired storage system ID to the **sr-probe** command to identify the storage pools available within the specified storage system

```
xe sr-probe type=cslg \
device-config:target=192.168.128.10 \ device-config:storageSystemId=HP__EVA__50014380013C0240
<?xml version="1.0" encoding="iso-8859-1"?>
<csl__storagePoolInfoList>
  <csl__storagePoolInfo>
    <displayName>Default Disk Group</displayName>
    <friendlyName>Default Disk Group</friendlyName>
    <storagePoolId>00010710B4080560B6AB0800008000000000000400</storagePoolId>
    <parentStoragePoolId></parentStoragePoolId>
    <storageSystemId>HP__EVA__50014380013C0240</storageSystemId>
    <sizeInMB>1957099</sizeInMB>
    <freeSpaceInMB>1273067</freeSpaceInMB>
    <isDefault>No</isDefault>
    <status>0</status>
    <provisioningOptions>
      <supportedRaidTypes>
        <raidType>RAID0</raidType>
        <raidType>RAID1</raidType>
        <raidType>RAID5</raidType>
      </supportedRaidTypes>
      <supportedNodeTypes>
        <nodeType>STORAGE_VOLUME</nodeType>
      </supportedNodeTypes>
      <supportedProvisioningTypes>
      </supportedProvisioningTypes>
    </provisioningOptions>
  </csl__storagePoolInfo>
</csl__storagePoolInfoList>
```

You can use `grep` to filter the `sr-probe` output to just the storage pool IDs

```
xe sr-probe type=cslg \
device-config:target=192.168.128.10 \
device-config:storageSystemId=HP__EVA__50014380013C0240 \
| grep storagePoolId
<storagePoolId>00010710B4080560B6AB0800008000000000000400</storagePoolId>
```

5. Create the SR specifying the desired storage system and storage pool IDs

```
xe sr-create type=cslg name-label=CSLG_EVA_1 shared=true \
device-config:target=192.168.128.10 \
device-config:storageSystemId=HP__EVA__50014380013C0240 \
device-config:storagePoolId=00010710B4080560B6AB0800008000000000000400
```

Managing Storage Repositories

This section covers various operations required in the ongoing management of Storage Repositories (SRs).

Destroying or forgetting a SR

You can destroy an SR, which actually deletes the contents of the SR from the physical media. Alternatively you can forget an SR, which allows you to re-attach the SR, for example, to another XenServer host, without removing any of the SR contents. In both cases, the PBD of the SR must first be unplugged. Forgetting an SR is the equivalent of the **SR Detach** operation within XenCenter.

1. Unplug the PBD to detach the SR from the corresponding XenServer host:

```
xe pbd-unplug uuid=<pbid_uuid>
```

2. To destroy the SR, which deletes both the SR and corresponding PBD from the XenServer host database and deletes the SR contents from the physical media:

```
xe sr-destroy uuid=<sr_uuid>
```

3. Or, to forget the SR, which removes the SR and corresponding PBD from the XenServer host database but leaves the actual SR contents intact on the physical media:

```
xe sr-forget uuid=<sr_uuid>
```

Note:

It might take some time for the software object corresponding to the SR to be garbage collected.

Introducing an SR

Introducing an SR that has been forgotten requires introducing an SR, creating a PBD, and manually plugging the PBD to the appropriate XenServer hosts to activate the SR.

The following example introduces a SR of type `lvmoiscsi`.

1. Probe the existing SR to determine its UUID:

```
xe sr-probe type=lvmoiscsi device-config:target=<192.168.1.10> \
device-config:targetIQN=<192.168.1.10:filer1> \
device-config:SCSIid=<14945540000000000000000000200000b70200000f000000>
```

2. Introduce the existing SR UUID returned from the **sr-probe** command. The UUID of the new SR is returned:

```
xe sr-introduce content-type=user name-label=<"Example Shared LVM over iSCSI SR">
shared=true uuid=<valid_sr_uuid> type=lvmoiscsi
```

3. Create a PBD to accompany the SR. The UUID of the new PBD is returned:

```
xe pbd-create type=lvmoiscsi host-uuid=<valid_uuid> sr-uuid=<valid_sr_uuid> \
device-config:target=<192.168.0.1> \
device-config:targetIQN=<192.168.1.10:filer1> \
device-config:SCSIid=<14945540000000000000000000200000b70200000f000000>
```

4. Plug the PBD to attach the SR:

```
xe pbd-plug uuid=<pbid_uuid>
```

5. Verify the status of the PBD plug. If successful the `currently-attached` property will be true:

```
xe pbd-list sr-uuid=<sr_uuid>
```

Note:

Steps 3 through 5 must be performed for each host in the resource pool, and can also be performed using the **Repair Storage Repository** function in XenCenter.

Resizing an SR

If you have resized the LUN on which a iSCSI or HBA SR is based, use the following procedures to reflect the size change in XenServer:

1. **iSCSI SRs** - unplug all PBDs on the host that reference LUNs on the same target. This is required to reset the iSCSI connection to the target, which in turn will allow the change in LUN size to be recognized when the PBDs are replugged.
2. **HBA SRs** - reboot the host.

Note:

In previous versions of XenServer explicit commands were required to resize the physical volume group of iSCSI and HBA SRs. These commands are now issued as part of the PBD plug operation and are no longer required.

Converting local Fibre Channel SRs to shared SRs

Use the xe CLI and the XenCenter **Repair Storage Repository** feature to convert a local FC SR to a shared FC SR:

1. Upgrade all hosts in the resource pool to XenServer 5.6.
2. Ensure all hosts in the pool have the SR's LUN zoned appropriately. See [the section called "Probing an SR"](#) for details on using the **sr-probe** command to verify the LUN is present on each host.
3. Convert the SR to shared:

```
xe sr-param-set shared=true uuid=<local_fc_sr>
```

4. Within XenCenter the SR is moved from the host level to the pool level, indicating that it is now shared. The SR will be marked with a red exclamation mark to show that it is not currently plugged on all hosts in the pool.
5. Select the SR and then select the **Storage > Repair Storage Repository** menu option.
6. Click **Repair** to create and plug a PBD for each host in the pool.

Moving Virtual Disk Images (VDIs) between SRs

The set of VDIs associated with a VM can be copied from one SR to another to accommodate maintenance requirements or tiered storage configurations. XenCenter provides the ability to copy a VM and all of its VDIs to the same or a different SR, and a combination of XenCenter and the xe CLI can be used to copy individual VDIs.

Copying all of a VM's VDIs to a different SR

The XenCenter **Copy VM** function creates copies of all VDIs for a selected VM on the same or a different SR. The source VM and VDIs are not affected by default. To move the VM to the selected SR rather than creating a copy, select the **Remove original VM** option in the **Copy Virtual Machine** dialog box.

1. Shutdown the VM.
2. Within XenCenter select the VM and then select the **VM > Copy VM** menu option.
3. Select the desired target SR.

Copying individual VDIs to a different SR

A combination of the xe CLI and XenCenter can be used to copy individual VDIs between SRs.

1. Shutdown the VM.
2. Use the xe CLI to identify the UUIDs of the VDIs to be moved. If the VM has a DVD drive its `vdi-uuid` will be listed as `<not in database>` and can be ignored.

```
xe vbd-list vm-uuid=<valid_vm_uuid>
```

Note:

The **vbd-list** command displays both the VBD and VDI UUIDs. Be sure to record the VDI UUIDs rather than the VBD UUIDs.

3. In XenCenter select the VM's **Storage** tab. For each VDI to be moved, select the VDI and click the **Detach** button. This step can also be done using the **vbd-destroy** command.

Note:

If you use the **vbd-destroy** command to detach the VDI UUIDs, be sure to first check if the VBD has the parameter `other-config:owner` set to `true`. If so, set it to `false`. Issuing the **vbd-destroy** command with `other-config:owner=true` will also destroy the associated VDI.

4. Use the **vdi-copy** command to copy each of the VM's VDIs to be moved to the desired SR.

```
xe vdi-copy uuid=<valid_vdi_uuid> sr-uuid=<valid_sr_uuid>
```

5. Within XenCenter select the VM's **Storage** tab. Click the **Attach** button and select the VDIs from the new SR. This step can also be done use the `vbd-create` command.
6. To delete the original VDIs, within XenCenter select the **Storage** tab of the original SR. The original VDIs will be listed with an empty value for the VM field and can be deleted with the Delete button.

Adjusting the disk IO scheduler

For general performance, the default disk scheduler `noop` is applied on all new SR types. The `noop` scheduler provides the fairest performance for competing VMs accessing the same device. To apply disk QoS (see [the section called "Virtual disk QoS settings"](#)) it is necessary to override the default setting and assign the `cfq` disk scheduler to the SR. The corresponding PBD must be unplugged and re-plugged for the scheduler parameter to take effect. The disk scheduler can be adjusted using the following command:

```
xe sr-param-set other-config:scheduler=noop|cfq|anticipatory|deadline \
uuid=<valid_sr_uuid>
```

Note:

This will not effect EqualLogic, NetApp or NFS storage.

Virtual disk QoS settings

Virtual disks have an optional I/O priority Quality of Service (QoS) setting. This setting can be applied to existing virtual disks using the xe CLI as described in this section.

In the shared SR case, where multiple hosts are accessing the same LUN, the QoS setting is applied to VBDs accessing the LUN from the same host. QoS is not applied across hosts in the pool.

Before configuring any QoS parameters for a VBD, ensure that the disk scheduler for the SR has been set appropriately. See [the section called "Adjusting the disk IO scheduler"](#) for details on how to adjust the scheduler. The scheduler parameter must be set to `cfq` on the SR for which the QoS is desired.

Note:



Remember to set the scheduler to `cfq` on the SR, and to ensure that the PBD has been re-plugged in order for the scheduler change to take effect.

The first parameter is `qos_algorithm_type`. This parameter needs to be set to the value **ionice**, which is the only type of QoS algorithm supported for virtual disks in this release.

The QoS parameters themselves are set with key/value pairs assigned to the `qos_algorithm_param` parameter. For virtual disks, `qos_algorithm_param` takes a `sched` key, and depending on the value, also requires a `class` key.

Possible values of `qos_algorithm_param:sched` are:

- `sched=rt` or `sched=real-time` sets the QoS scheduling parameter to real time priority, which requires a `class` parameter to set a value
- `sched=idle` sets the QoS scheduling parameter to idle priority, which requires no `class` parameter to set any value
- `sched=<anything>` sets the QoS scheduling parameter to best effort priority, which requires a `class` parameter to set a value

The possible values for `class` are:

- One of the following keywords: `highest`, `high`, `normal`, `low`, `lowest`
- an integer between 0 and 7, where 7 is the highest priority and 0 is the lowest, so that, for example, I/O requests with a priority of 5, will be given priority over I/O requests with a priority of 2.

To enable the disk QoS settings, you also need to set the `other-config:scheduler` to `cfq` and replug PBDs for the storage in question.

For example, the following CLI commands set the virtual disk's VBD to use real time priority 5:

```
xe vbd-param-set uuid=<vbd_uuid> qos_algorithm_type=ionice
xe vbd-param-set uuid=<vbd_uuid> qos_algorithm_params:sched=rt
xe vbd-param-set uuid=<vbd_uuid> qos_algorithm_params:class=5
xe sr-param-set uuid=<sr_uuid> other-config:scheduler=cfq
xe pbd-plug uuid=<pbd_uuid>
```

Configuring VM memory

When a VM is first created, it is allocated a fixed amount of memory. To improve the utilization of physical memory in your XenServer environment, you can use Dynamic Memory Control (DMC), a memory management feature that enables dynamic reallocation of memory between VMs.

XenCenter provides a graphical display of memory usage in its **Memory** tab. This is described in the *XenCenter Help*.

In previous editions of XenServer adjusting virtual memory on VMs required a restart to add or remove memory and an interruption to users' service.

Dynamic Memory Control (DMC) provides the following benefits:

- Memory can be added or removed without restart thus providing a more seamless experience to the user.
- When servers are full, DMC allows you to start more VMs on these servers, reducing the amount of memory allocated to the running VMs proportionally.

What is Dynamic Memory Control (DMC)?

XenServer DMC (sometimes known as "dynamic memory optimization", "memory overcommit" or "memory ballooning") works by automatically adjusting the memory of running VMs, keeping the amount of memory allocated to each VM between specified minimum and maximum memory values, guaranteeing performance and permitting greater density of VMs per server. Without DMC, when a server is full, starting further VMs will fail with "out of memory" errors: to reduce the existing VM memory allocation and make room for more VMs you must edit each VM's memory allocation and then reboot the VM. With DMC enabled, even when the server is full, XenServer will attempt to reclaim memory by automatically reducing the current memory allocation of running VMs within their defined memory ranges.

Without DMC, when a server is full, starting further VMs will fail with "out of memory" errors: to reduce the existing VM memory allocation and make room for more VMs you must edit each VM's memory allocation and then reboot the VM. With DMC enabled, even when the server is full, XenServer will attempt to reclaim memory by automatically reducing the current memory allocation of running VMs within their defined memory ranges.

Note:

Dynamic Memory Control is only available for XenServer Advanced or higher editions. To learn more about XenServer Advanced or higher editions and to find out how to upgrade, visit the Citrix website [here](#)

The concept of dynamic range

For each VM the administrator can set a dynamic memory range – this is the range within which memory can be added/removed from the VM without requiring a reboot. When a VM is running the administrator can adjust the dynamic range. XenServer always guarantees to keep the amount of memory allocated to the VM within the dynamic range; therefore adjusting it while the VM is running may cause XenServer to adjust the amount of memory allocated to the VM. (The most extreme case is where the administrator sets the dynamic min/max to the same value, thus forcing XenServer to ensure that this amount of memory is allocated to the VM.) If new VMs are required to start on "full" servers, running VMs have their memory 'squeezed' to start new ones. The required extra memory is obtained by squeezing the existing running VMs proportionally within their pre-defined dynamic ranges

DMC allows you to configure dynamic minimum and maximum memory levels – creating a Dynamic Memory Range (DMR) that the VM will operate in.



- Dynamic Minimum Memory: A lower memory limit that you assign to the VM.
- Dynamic Higher Limit: An upper memory limit that you assign to the VM.

For example, if the Dynamic Minimum Memory was set at 512 MB and the Dynamic Maximum Memory was set at 1024 MB this would give the VM a Dynamic Memory Range (DMR) of 512 - 1024 MB, within which, it would operate. With DMC, XenServer *guarantees* at all times to assign each VM memory within its specified DMR.

The concept of static range

Many Operating Systems that XenServer supports do not fully 'understand' the notion of dynamically adding or removing memory. As a result, XenServer must declare the maximum amount of memory that a VM will ever be asked to consume at the time that it boots. (This allows the guest operating system to size its page tables and other memory management structures accordingly.) This introduces the concept of a static memory range within XenServer. The static memory range cannot be adjusted while the VM is running. For a particular boot, the dynamic range is constrained such as to be always contained within this static range. Note that the static minimum (the lower bound of the static range) is there to protect the administrator and is set to the lowest amount of memory that the OS can run with on XenServer.

Note:

Citrix advises not to change the static minimum level as this is set at the supported level per operating system – refer to the memory constraints table for more details.

By setting a static maximum level, higher than a dynamic max, means that in the future, if you need to allocate more memory to a VM, you can do so without requiring a reboot.

DMC Behaviour

Automatic VM squeezing

- If DMC is not enabled, when hosts are full, new VM starts fail with 'out of memory' errors.
- If DMC is enabled, even when hosts are full, XenServer will attempt to reclaim memory (by reducing the memory allocation of running VMs within their defined dynamic ranges). In this way running VMs are squeezed proportionally at the same distance between the dynamic minimum and dynamic maximum for all VMs on the host

When DMC is enabled

- When the host's memory is plentiful - All running VMs will receive their Dynamic Maximum Memory level
- When the host's memory is scarce - All running VMs will receive their Dynamic Minimum Memory level.

When you are configuring DMC, remember that allocating only a small amount of memory to a VM can negatively impact it. For example, allocating too little memory:

- Using Dynamic Memory Control to reduce the amount of physical memory available to a VM may cause it to boot slowly. Likewise, if you allocate too little memory to a VM, it may start extremely slowly.
- Setting the dynamic memory minimum for a VM too low may result in poor performance or stability problems when the VM is starting.

How does DMC Work?

Using DMC, it is possible to operate a guest virtual machine in one of two modes:

1. Target Mode: The administrator specifies a memory target for the guest. XenServer adjusts the guest's memory allocation to meet the target. Specifying a target is particularly useful in virtual server environments, and in any situation where you know exactly how much memory you want a guest to use. XenServer will adjust the guest's memory allocation to meet the target you specify.
2. Dynamic Range Mode: The administrator specifies a dynamic memory range for the guest; XenServer chooses a target from within the range and adjusts the guest's memory allocation to meet the target. Specifying a dynamic range is particularly useful in virtual desktop environments, and in any situation where you want XenServer to repartition host memory dynamically in response to changing numbers of guests, or changing host memory pressure. XenServer chooses a target from within the range and adjusts the guest's memory allocation to meet the target.

Note:

It is possible to change between target mode and dynamic range mode at any time for any running guest. Simply specify a new target, or a new dynamic range, and XenServer takes care of the rest.

Memory constraints

XenServer allows administrators to use all memory control operations with any guest operating system. However, XenServer enforces the following memory property ordering constraint for all guests:

$$0 \leq \text{memory-static-min} \leq \text{memory-dynamic-min} \leq \text{memory-dynamic-max} \leq \text{memory-static-max}$$

XenServer allows administrators to change guest memory properties to any values that satisfy this constraint, subject to validation checks. However, in addition to the above constraint, Citrix supports only certain guest memory configurations for each supported operating system. See below for further details.

Supported operating systems

Citrix supports only certain guest memory configurations. The range of supported configurations depends on the guest operating system in use. XenServer does not prevent administrators from configuring guests to exceed the supported limit. However, customers are strongly advised to keep memory properties within the supported limits to avoid performance or stability problems.

Operating System			Supported Memory Limits		
Family	Version	Architectures	Dynamic Minimum	Dynamic Maximum	Additional Constraints
Microsoft Windows	XP (SP2, SP3)	x86	$\geq 256 \text{ MB}$	$\leq 4 \text{ GB}$	Dynamic Minimum $\geq \frac{1}{4}$ Static Maximum for all supported operating systems
	Server 2003	x86 x64	$\geq 256 \text{ MB}$	$\leq 32 \text{ GB}$	
	Server 2008	x86 x64	$\geq 512 \text{ MB}$	$\leq 32 \text{ GB}$	
	Server 2008 R2	x86 x64	$\geq 512 \text{ MB}$	$\leq 32 \text{ GB}$	
	Vista	x86	$\geq 1 \text{ GB}$	$\leq 4 \text{ GB}$	
	7	x86	$\geq 1 \text{ GB}$	$\leq 4 \text{ GB}$	

Operating System			Supported Memory Limits	
		x64	≥ 2 GB	≤ 32 GB
CentOS Linux	4.5 - 4.8	x86	≥ 256 MB	≤ 16 GB
	5.0 - 5.4	x86 x64	≥ 512 MB	≤ 16 GB
RedHat Enterprise Linux	4.5 - 4.8	x86	≥ 256 MB	≤ 16 GB
	5.0 - 5.4	x86 x64	≥ 512 MB	≤ 16 GB
Oracle Enterprise Linux	5.0 - 5.4	x86 x64	≥ 512 MB	≤ 16 GB
SUSE Enterprise Linux	10 (SP1, SP2)	x86 x64	≥ 512 MB	≤ 32 GB
	11	x86 x64	≥ 512 MB	≤ 32 GB
Debian GNU/Linux	Lenny	x86	≥ 128 MB	≤ 32 GB

Warning:

When configuring guest memory, please be careful NOT to exceed the maximum amount of physical memory addressable by your operating system. Setting a memory maximum that is greater than the operating system supported limit, may lead to stability problems within your guest.

xe CLI commands

Display the static memory properties of a VM

1. Find the uuid of the required VM:

```
xe vm-list
```

2. Note the uuid, and then run the command **param-name=memory-static**

```
xe vm-param-get uuid=<uuid> param-name=memory-static-{min,max}
```

For example, the following displays the static maximum memory properties for the VM with the uuid beginning ec77:

```
xe vm-param-get uuid= \
ec77a893-bff2-aa5c-7ef2-9c3acf0f83c0 \
param-name=memory-static-max;
268435456
```

This shows that the static maximum memory for this VM is 268435456 bytes (256MB).

Display the dynamic memory properties of a VM

To display the dynamic memory properties, follow the procedure as above but use the command **param-name=memory-dynamic**:

1. Find the uuid of the required VM:

```
xe vm-list
```

2. Note the uuid, and then run the command **param-name=memory-dynamic**:

```
xe vm-param-get uuid=<uuid> param-name=memory-dynamic-{min,max}
```

For example, the following displays the dynamic maximum memory properties for the VM with uuid beginning ec77

```
xe vm-param-get uuid= \
  ec77a893-bff2-aa5c-7ef2-9c3acf0f83c0 \
  param-name=memory-dynamic-max;
134217728
```

This shows that the static maximum memory for this VM is 134217728 bytes (128MB).

Updating memory properties

Warning:

It is essential that you use the correct ordering when setting the static/dynamic minimum/maximum parameters. In addition you must not invalidate the following constraint:

$$0 \leq \text{memory-static-min} \leq \text{memory-dynamic-min} \leq \text{memory-dynamic-max} \leq \text{memory-static-max}$$

Update the static memory range of a virtual machine:

```
xe vm-memory-static-range-set uuid=<uuid> min=<value>max=<value>
```

Update the dynamic memory range of a virtual machine:

```
xe vm-memory-dynamic-range-set \
  uuid=<uuid> min=<value> \
  max=<value>
```

Specifying a target is particularly useful in virtual server environments, and in any situation where you know exactly how much memory you want a guest to use. XenServer will adjust the guest's memory allocation to meet the target you specify. For example:

```
xe vm-target-set
```

Update all memory limits (static and dynamic) of a virtual machine:

```
xe vm-memory-limits-set \
  uuid=<uuid> \
  static-min=<value> \
  dynamic-min=<value> \
  dynamic-max=<value> static-max=<value>
```

Note:

- To allocate a specific amount memory to a VM that won't change, set the Dynamic Maximum and Dynamic Minimum to the same value.
- You cannot increase the dynamic memory of a VM beyond the static maximum.
- To alter the static maximum of a VM – you will need to suspend or shut down the VM.



Update individual memory properties

Warning:

Citrix advises not to change the static minimum level as this is set at the supported level per operating system – refer to the memory constraints table for more details.

Update the dynamic memory properties of a VM.

1. Find the uuid of the required VM:

```
xe vm-list
```

2. Note the uuid, and then use the command **memory-dynamic-{min,max}=<value>**

```
xe vm-param-set uuid=<uuid>memory-dynamic-{min,max}=<value>
```

The following example changes the dynamic maximum to 128MB:

```
xe vm-param-set uuid=ec77a893-bff2-aa5c-7ef2-9c3acf0f83c0 memory-dynamic-max=128MiB
```

Upgrade issues

After upgrading from Citrix XenServer 5.5, XenServer sets all VMs memory so that the dynamic minimum is equal to the dynamic maximum.

Workload Balancing interaction

If Workload Balancing (WLB) is enabled, XenServer defers decisions about host selection to the workload balancing server. If WLB is disabled, or if the WLB server has failed or is unavailable, XenServer will use its internal algorithm to make decisions regarding host selection.

Networking

This chapter discusses how physical network interface cards (NICs) in XenServer hosts are used to enable networking within Virtual Machines (VMs). XenServer supports up to 16 physical network interfaces (or up to 16 of bonded network interfaces) per XenServer host and up to 7 virtual network interfaces per VM.

Note:

XenServer provides automated configuration and management of NICs using the xe command line interface (CLI). Unlike previous XenServer versions, the host networking configuration files should not be edited directly in most cases; where a CLI command is available, do not edit the underlying files.

If you are already familiar with XenServer networking concepts, you may want to skip ahead to one of the following sections:

- For procedures on how to create networks for standalone XenServer hosts, see [the section called “Creating networks in a standalone server”](#).
- For procedures on how to create networks for XenServer hosts that are configured in a resource pool, see [the section called “Creating networks in resource pools”](#).
- For procedures on how to create VLANs for XenServer hosts, either standalone or part of a resource pool, see [the section called “Creating VLANs”](#).
- For procedures on how to create bonds for standalone XenServer hosts, see [the section called “Creating NIC bonds on a standalone host”](#).
- For procedures on how to create bonds for XenServer hosts that are configured in a resource pool, see [the section called “Creating NIC bonds in resource pools”](#).

XenServer networking overview

This section describes the general concepts of networking in the XenServer environment.

One network is created for each physical network interface card during XenServer installation. When you add a server to a resource pool, these default networks are merged so that all physical NICs with the same device name are attached to the same network.

Typically you would only add a new network if you wished to create an internal network, set up a new VLAN using an existing NIC, or create a NIC bond.

You can configure three different types of physical (server) networks in XenServer:

- **Internal networks** have no association to a physical network interface, and can be used to provide connectivity only between the virtual machines on a given server, with no connection to the outside world.
- **External networks** have an association with a physical network interface and provide a bridge between a virtual machine and the physical network interface connected to the network, enabling a virtual machine to connect to resources available through the server’s physical network interface card.
- **Bonded networks** create a bond between two NICs to create a single, high-performing channel between the virtual machine and the network.

Note:

Some networking options have different behaviors when used with standalone XenServer hosts compared to resource pools. This chapter contains sections on general information that applies to both standalone hosts and pools, followed by specific information and procedures for each.

Network objects

There are three types of server-side software objects which represent networking entities. These objects are:

- A *PIF*, which represents a physical network interface on a XenServer host. PIF objects have a name and description, a globally unique UUID, the parameters of the NIC that they represent, and the network and server they are connected to.
- A *VIF*, which represents a virtual interface on a Virtual Machine. VIF objects have a name and description, a globally unique UUID, and the network and VM they are connected to.
- A *network*, which is a virtual Ethernet switch on a XenServer host. Network objects have a name and description, a globally unique UUID, and the collection of VIFs and PIFs connected to them.

Both XenCenter and the xe CLI allow configuration of networking options, control over which NIC is used for management operations, and creation of advanced networking features such as virtual local area networks (VLANs) and NIC bonds.

Networks

Each XenServer host has one or more networks, which are virtual Ethernet switches. Networks without an association to a PIF are considered *internal*, and can be used to provide connectivity only between VMs on a given XenServer host, with no connection to the outside world. Networks with a PIF association are considered *external*, and provide a bridge between VIFs and the PIF connected to the network, enabling connectivity to resources available through the PIF's NIC.

VLANs

Virtual Local Area Networks (VLANs), as defined by the IEEE 802.1Q standard, allow a single physical network to support multiple logical networks. XenServer hosts can work with VLANs in multiple ways.

Note:

All supported VLAN configurations are equally applicable to pools and standalone hosts, and bonded and non-bonded configurations.

Using VLANs with host management interfaces

Switch ports configured to perform 802.1Q VLAN tagging/untagging, commonly referred to as ports with a *native VLAN* or as *access mode* ports, can be used with XenServer management interfaces to place management traffic on a desired VLAN. In this case the XenServer host is unaware of any VLAN configuration.

XenServer management interfaces cannot be assigned to a XenServer VLAN via a trunk port.

Using VLANs with virtual machines

Switch ports configured as 802.1Q VLAN trunk ports can be used in combination with the XenServer VLAN features to connect guest virtual network interfaces (VIFs) to specific VLANs. In this case the XenServer host performs the VLAN tagging/untagging functions for the guest, which is unaware of any VLAN configuration.

XenServer VLANs are represented by additional PIF objects representing VLAN interfaces corresponding to a specified VLAN tag. XenServer networks can then be connected to the PIF representing the physical NIC to see all traffic on the NIC, or to a PIF representing a VLAN to see only the traffic with the specified VLAN tag.

For procedures on how to create VLANs for XenServer hosts, either standalone or part of a resource pool, see [the section called “Creating VLANs”](#).

Using VLANs with dedicated storage NICs

Dedicated storage NICs can be configured to use native VLAN / access mode ports as described above for management interfaces, or with trunk ports and XenServer VLANs as described above for virtual machines. To configure dedicated storage NICs, see [the section called “Configuring a dedicated storage NIC”](#).

Combining management interfaces and guest VLANs on a single host NIC

A single switch port can be configured with both trunk and native VLANs, allowing one host NIC to be used for a management interface (on the native VLAN) and for connecting guest VIFs to specific VLAN IDs.

NIC bonds

NIC bonds can improve XenServer host resiliency by using two physical NICs as if they were one. If one NIC within the bond fails the host's network traffic will automatically be routed over the second NIC. NIC bonds work in an active/active mode, with traffic balanced between the bonded NICs.

XenServer NIC bonds completely subsume the underlying physical devices (PIFs). In order to activate a bond the underlying PIFs must not be in use, either as the management interface for the host or by running VMs with VIFs attached to the networks associated with the PIFs.

XenServer NIC bonds are represented by additional PIFs. The bond PIF can then be connected to a XenServer network to allow VM traffic and host management functions to occur over the bonded NIC. The exact steps to use to create a NIC bond depend on the number of NICs in your host, and whether the management interface of the host is assigned to a PIF to be used in the bond.

XenServer supports Source Level Balancing (SLB) NIC bonding. SLB bonding:

- is an active/active mode, but only supports load-balancing of VM traffic across the physical NICs
- provides fail-over support for all other traffic types
- does not require switch support for Etherchannel or 802.3ad (LACP)
- load balances traffic between multiple interfaces at VM granularity by sending traffic through different interfaces based on the source MAC address of the packet
- is derived from the open source ALB mode and reuses the ALB capability to dynamically re-balance load across interfaces

Any given VIF will only use one of the links in the bond at a time. At startup no guarantees are made about the affinity of a given VIF to a link in the bond. However, for VIFs with high throughput, periodic rebalancing ensures that the load on the links is approximately equal.

API Management traffic can be assigned to a XenServer bond interface and will be automatically load-balanced across the physical NICs.

XenServer bonded PIFs do not require IP configuration for the bond when used for guest traffic. This is because the bond operates at Layer 2 of the OSI, the data link layer, and no IP addressing is used at this layer. When used for non-guest traffic (to connect to it with XenCenter for management, or to connect to shared network storage), one IP configuration is required per bond. (Incidentally, this is true of unbonded PIFs as well, and is unchanged from XenServer 4.1.0.)

Gratuitous ARP packets are sent when assignment of traffic changes from one interface to another as a result of fail-over.

Re-balancing is provided by the existing ALB re-balance capabilities: the number of bytes going over each slave (interface) is tracked over a given period. When a packet is to be sent that contains a new source MAC address it is assigned to the slave interface with the lowest utilization. Traffic is re-balanced every 10 seconds.

Note:

Bonding is set up with an Up Delay of 31000ms and a Down Delay of 200ms. The seemingly long Up Delay is purposeful because of the time taken by some switches to actually start routing traffic. Without it, when a link comes back after failing, the bond might rebalance traffic onto it before the switch is ready to pass traffic. If you want to move both connections to a different switch, move one, then wait 31 seconds for it to be used again before moving the other.

Initial networking configuration

The XenServer host networking configuration is specified during initial host installation. Options such as IP address configuration (DHCP/static), the NIC used as the management interface, and hostname are set based on the values provided during installation.

When a XenServer host has a single NIC, the follow configuration is present after installation:

- a single PIF is created corresponding to the host's single NIC
- the PIF is configured with the IP addressing options specified during installation and to enable management of the host
- the PIF is set for use in host management operations
- a single network, network 0, is created
- network 0 is connected to the PIF to enable external connectivity to VMs

When a host has multiple NICs the configuration present after installation depends on which NIC is selected for management operations during installation:

- PIFs are created for each NIC in the host
- the PIF of the NIC selected for use as the management interface is configured with the IP addressing options specified during installation
- a network is created for each PIF ("network 0", "network 1", etc.)
- each network is connected to one PIF
- the IP addressing options of all other PIFs are left unconfigured

In both cases the resulting networking configuration allows connection to the XenServer host by XenCenter, the xe CLI, and any other management software running on separate machines via the IP address of the management interface. The configuration also provides external networking for VMs created on the host.

The PIF used for management operations is the only PIF ever configured with an IP address. External networking for VMs is achieved by bridging PIFs to VIFs using the network object which acts as a virtual Ethernet switch.

The steps required for networking features such as VLANs, NIC bonds, and dedicating a NIC to storage traffic are covered in the following sections.

Managing networking configuration

Some of the network configuration procedures in this section differ depending on whether you are configuring a stand-alone server or a server that is part of a resource pool.

Creating networks in a standalone server

Because external networks are created for each PIF during host installation, creating additional networks is typically only required to:



- use an internal network
- support advanced operations such as VLANs or NIC bonding

To add or remove networks using XenCenter, refer to the XenCenter online Help.

To add a new network using the CLI

1. Open the XenServer host text console.
2. Create the network with the `network-create` command, which returns the UUID of the newly created network:

```
xe network-create name-label=<mynetwork>
```

At this point the network is not connected to a PIF and therefore is internal.

Creating networks in resource pools

All XenServer hosts in a resource pool should have the same number of physical network interface cards (NICs), although this requirement is not strictly enforced when a XenServer host is joined to a pool.

Having the same physical networking configuration for XenServer hosts within a pool is important because all hosts in a pool share a common set of XenServer networks. PIFs on the individual hosts are connected to pool-wide networks based on device name. For example, all XenServer hosts in a pool with an `eth0` NIC will have a corresponding PIF plugged into the pool-wide `Network 0` network. The same will be true for hosts with `eth1` NICs and `Network 1`, as well as other NICs present in at least one XenServer host in the pool.

If one XenServer host has a different number of NICs than other hosts in the pool, complications can arise because not all pool networks will be valid for all pool hosts. For example, if hosts `host1` and `host2` are in the same pool and `host1` has four NICs while `host2` only has two, only the networks connected to PIFs corresponding to `eth0` and `eth1` will be valid on `host2`. VMs on `host1` with VIFs connected to networks corresponding to `eth2` and `eth3` will not be able to migrate to host `host2`.

All NICs of all XenServer hosts within a resource pool must be configured with the same MTU size.

Creating VLANs

For servers in a resource pool, you can use the **pool-vlan-create** command. This command creates the VLAN and automatically creates and plugs in the required PIFs on the hosts in the pool. See [the section called “pool-vlan-create”](#) for more information.

To connect a network to an external VLAN using the CLI

1. Open the XenServer host text console.
2. Create a new network for use with the VLAN. The UUID of the new network is returned:

```
xe network-create name-label=network5
```

3. Use the **pif-list** command to find the UUID of the PIF corresponding to the physical NIC supporting the desired VLAN tag. The UUIDs and device names of all PIFs are returned, including any existing VLANs:

```
xe pif-list
```

4. Create a VLAN object specifying the desired physical PIF and VLAN tag on all VMs to be connected to the new VLAN. A new PIF will be created and plugged into the specified network. The UUID of the new PIF object is returned.

```
xe vlan-create network-uuid=<network_uuid> pif-uuid=<pif_uuid> vlan=5
```

5. Attach VM VIFs to the new network. See [the section called “Creating networks in a standalone server”](#) for more details.

Creating NIC bonds on a standalone host

Citrix recommends using XenCenter to create NIC bonds. For details, refer to the XenCenter help.

This section describes how to use the xe CLI to create bonded NIC interfaces on a standalone XenServer host. See [the section called “Creating NIC bonds in resource pools”](#) for details on using the xe CLI to create NIC bonds on XenServer hosts that comprise a resource pool.

Creating a NIC bond on a dual-NIC host

Creating a bond on a dual-NIC host implies that the PIF/NIC currently in use as the management interface for the host will be subsumed by the bond. The additional steps required to move the management interface to the bond PIF are included.

Bonding two NICs together

1. Use XenCenter or the **vm-shutdown** command to shut down all VMs on the host, thereby forcing all VIFs to be unplugged from their current networks. The existing VIFs will be invalid after the bond is enabled.

```
xe vm-shutdown uuid=<vm_uuid>
```

2. Use the **network-create** command to create a new network for use with the bonded NIC. The UUID of the new network is returned:

```
xe network-create name-label=<bond0>
```

3. Use the **pif-list** command to determine the UUIDs of the PIFs to use in the bond:

```
xe pif-list
```

4. Use the **bond-create** command to create the bond; separated by commas, specify the newly created network UUID and the UUIDs of the PIFs to be bonded. The UUID for the bond is returned:

```
xe bond-create network-uuid=<network_uuid> pif-uuids=<pif_uuid_1>,<pif_uuid_2>
```

Note:

See [the section called “Controlling the MAC address of the bond”](#) for details on controlling the MAC address used for the bond PIF.

5. Use the **pif-list** command to determine the UUID of the new bond PIF:

```
xe pif-list device=<bond0>
```

6. Use the **pif-reconfigure-ip** command to configure the desired management interface IP address settings for the bond PIF. See [Appendix A, Command line interface](#) for more detail on the options available for the pif-reconfigure-ip command.

```
xe pif-reconfigure-ip uuid=<bond_pif_uuid> mode=DHCP
```

7. Use the **host-management-reconfigure** command to move the management interface from the existing physical PIF to the bond PIF. This step will activate the bond:

```
xe host-management-reconfigure pif-uuid=<bond_pif_uuid>
```

8. Use the **pif-reconfigure-ip** command to remove the IP address configuration from the non-bonded PIF previously used for the management interface. This step is not strictly necessary but might help reduce confusion when reviewing the host networking configuration.

```
xe pif-reconfigure-ip uuid=<old_management_pif_uuid> mode=None
```

9. Move existing VMs to the bond network using the **vif-destroy** and **vif-create** commands. This step can also be completed using XenCenter by editing the VM configuration and connecting the existing VIFs of a VM to the bond network.
10. Restart the VMs shut down in step 1.

Controlling the MAC address of the bond

Creating a bond on a dual-NIC host implies that the PIF/NIC currently in use as the management interface for the host will be subsumed by the bond. If DHCP is used to supply IP addresses to the host in most cases the MAC address of the bond should be the same as the PIF/NIC currently in use, allowing the IP address of the host received from DHCP to remain unchanged.

The MAC address of the bond can be changed from PIF/NIC currently in use for the management interface, but doing so will cause existing network sessions to the host to be dropped when the bond is enabled and the MAC/IP address in use changes.

The MAC address to be used for a bond can be controlled in two ways:

- an optional *mac* parameter can be specified in the **bond-create** command. Using this parameter, the bond MAC address can be set to any arbitrary address.
- If the *mac* parameter is not specified, the MAC address of the first PIF listed in the *pif-uuids* parameter is used for the bond.

Reverting NIC bonds

If reverting a XenServer host to a non-bonded configuration, be aware of the following requirements:

- As when creating a bond, all VMs with VIFs on the bond must be shut down prior to destroying the bond. After reverting to a non-bonded configuration, reconnect the VIFs to an appropriate network.
- Move the management interface to another PIF using the **pif-reconfigure-ip** and **host-management-reconfigure** commands prior to issuing the **bond-destroy** command, otherwise connections to the host (including XenCenter) will be dropped.

Creating NIC bonds in resource pools

Whenever possible, create NIC bonds as part of initial resource pool creation prior to joining additional hosts to the pool or creating VMs. Doing so allows the bond configuration to be automatically replicated to hosts as they are joined to the pool and reduces the number of steps required. Adding a NIC bond to an existing pool requires creating the bond configuration manually on the master and each of the members of the pool. Adding a NIC bond to an existing pool after VMs have been installed is also a disruptive operation, as all VMs in the pool must be shut down.

Citrix recommends using XenCenter to create NIC bonds. For details, refer to the XenCenter help.

This section describes using the xe CLI to create bonded NIC interfaces on XenServer hosts that comprise a resource pool. See [the section called “Creating a NIC bond on a dual-NIC host”](#) for details on using the xe CLI to create NIC bonds on a standalone XenServer host.

Warning:

Do not attempt to create network bonds while HA is enabled. The process of bond creation will disturb the in-progress HA heartbeating and cause hosts to self-fence (shut themselves down); subsequently they will likely fail to reboot properly and will need the **host-emergency-ha-disable** command to recover.

Adding NIC bonds to new resource pools

1. Select the host you want to be the master. The master host belongs to an unnamed pool by default. To create a resource pool with the CLI, rename the existing nameless pool:

```
xe pool-param-set name-label=<"New Pool"> uuid=<pool_uuid>
```

2. Create the NIC bond on the master as follows:

- a. Use the **network-create** command to create a new pool-wide network for use with the bonded NICs. The UUID of the new network is returned.

```
xe network-create name-label=<network_name>
```

- b. Use the **pif-list** command to determine the UUIDs of the PIFs to use in the bond:

```
xe pif-list
```

- c. Use the **bond-create** command to create the bond, specifying the network UUID created in step a and the UUIDs of the PIFs to be bonded, separated by commas. The UUID for the bond is returned:

```
xe bond-create network-uuid=<network_uuid> pif-uuids=<pif_uuid_1>,<pif_uuid_2>
```

Note:

See [the section called "Controlling the MAC address of the bond"](#) for details on controlling the MAC address used for the bond PIF.

- d. Use the **pif-list** command to determine the UUID of the new bond PIF:

```
xe pif-list network-uuid=<network_uuid>
```

- e. Use the **pif-reconfigure-ip** command to configure the desired management interface IP address settings for the bond PIF. See [Appendix A, Command line interface](#), for more detail on the options available for the **pif-reconfigure-ip** command.

```
xe pif-reconfigure-ip uuid=<bond_pif_uuid> mode=DHCP
```

- f. Use the **host-management-reconfigure** command to move the management interface from the existing physical PIF to the bond PIF. This step will activate the bond:

```
xe host-management-reconfigure pif-uuid=<bond_pif_uuid>
```

- g. Use the **pif-reconfigure-ip** command to remove the IP address configuration from the non-bonded PIF previously used for the management interface. This step is not strictly necessary but might help reduce confusion when reviewing the host networking configuration.

```
xe pif-reconfigure-ip uuid=<old_management_pif_uuid> mode=None
```

3. Open a console on a host that you want to join to the pool and run the command:

```
xe pool-join master-address=<host1> master-username=root master-password=<password>
```

The network and bond information is automatically replicated to the new host. However, the management interface is not automatically moved from the host NIC to the bonded NIC. Move the management interface on the host to enable the bond as follows:

- a. Use the **host-list** command to find the UUID of the host being configured:

```
xe host-list
```

- b. Use the **pif-list** command to determine the UUID of bond PIF on the new host. Include the *host-uuid* parameter to list only the PIFs on the host being configured:

```
xe pif-list network-name-label=<network_name> host-uuid=<host_uuid>
```

- c. Use the **pif-reconfigure-ip** command to configure the desired management interface IP address settings for the bond PIF. See [Appendix A, Command line interface](#), for more detail on the options available for the **pif-reconfigure-ip** command. *This command must be run directly on the host:*

```
xe pif-reconfigure-ip uuid=<bond_pif_uuid> mode=DHCP
```

- d. Use the **host-management-reconfigure** command to move the management interface from the existing physical PIF to the bond PIF. This step activates the bond. *This command must be run directly on the host:*

```
xe host-management-reconfigure pif-uuid=<bond_pif_uuid>
```

- e. Use the **pif-reconfigure-ip** command to remove the IP address configuration from the non-bonded PIF previously used for the management interface. This step is not strictly necessary but may help reduce confusion when reviewing the host networking configuration. *This command must be run directly on the host server:*

```
xe pif-reconfigure-ip uuid=<old_mgmt_pif_uuid> mode=None
```

4. For each additional host you want to join to the pool, repeat steps 3 and 4 to move the management interface on the host and to enable the bond.

Adding NIC bonds to an existing pool

Warning:

Do not attempt to create network bonds while HA is enabled. The process of bond creation disturbs the in-progress HA heartbeating and causes hosts to self-fence (shut themselves down); subsequently they will likely fail to reboot properly and you will need to run the **host-emergency-ha-disable** command to recover them.

Note:

If you are not using XenCenter for NIC bonding, the quickest way to create pool-wide NIC bonds is to create the bond on the master, and then restart the other pool members. Alternately you can use the **service xapi restart** command. This causes the bond and VLAN settings on the master to be inherited by each host. The management interface of each host must, however, be manually reconfigured.

When adding a NIC bond to an existing pool, the bond must be manually created on each host in the pool. The steps below can be used to add NIC bonds on both the pool master and other hosts with the following requirements:

1. All VMs in the pool must be shut down
2. Add the bond to the pool master first, and then to other hosts.
3. The **bond-create**, **host-management-reconfigure** and **host-management-disable** commands affect the host on which they are run and so are not suitable for use on one host in a pool to change the configuration of another. Run these commands directly on the console of the host to be affected.

To add NIC bonds to the pool master and other hosts

1. Use the **network-create** command to create a new pool-wide network for use with the bonded NICs. This step should only be performed once per pool. The UUID of the new network is returned.

```
xe network-create name-label=<bond0>
```

2. Use XenCenter or the **vm-shutdown** command to shut down all VMs in the host pool to force all existing VIFs to be unplugged from their current networks. The existing VIFs will be invalid after the bond is enabled.

```
xe vm-shutdown uuid=<vm_uuid>
```

3. Use the **host-list** command to find the UUID of the host being configured:

```
xe host-list
```

4. Use the **pif-list** command to determine the UUIDs of the PIFs to use in the bond. Include the *host-uuid* parameter to list only the PIFs on the host being configured:

```
xe pif-list host-uuid=<host_uuid>
```

5. Use the **bond-create** command to create the bond, specifying the network UUID created in step 1 and the UUIDs of the PIFs to be bonded, separated by commas. The UUID for the bond is returned.

```
xe bond-create network-uuid=<network_uuid> pif-uuids=<pif_uuid_1>,<pif_uuid_2>
```

Note:

See [the section called "Controlling the MAC address of the bond"](#) for details on controlling the MAC address used for the bond PIF.

6. Use the **pif-list** command to determine the UUID of the new bond PIF. Include the *host-uuid* parameter to list only the PIFs on the host being configured:

```
xe pif-list device=bond0 host-uuid=<host_uuid>
```

7. Use the **pif-reconfigure-ip** command to configure the desired management interface IP address settings for the bond PIF. See [Appendix A, Command line interface](#) for more detail on the options available for the **pif-reconfigure-ip** command. *This command must be run directly on the host:*

```
xe pif-reconfigure-ip uuid=<bond_pif_uuid> mode=DHCP
```

8. Use the **host-management-reconfigure** command to move the management interface from the existing physical PIF to the bond PIF. This step will activate the bond. *This command must be run directly on the host:*

```
xe host-management-reconfigure pif-uuid=<bond_pif_uuid>
```

9. Use the **pif-reconfigure-ip** command to remove the IP address configuration from the non-bonded PIF previously used for the management interface. This step is not strictly necessary, but might help reduce confusion when reviewing the host networking configuration. *This command must be run directly on the host:*

```
xe pif-reconfigure-ip uuid=<old_management_pif_uuid> mode=None
```

10. Move existing VMs to the bond network using the **vif-destroy** and **vif-create** commands. This step can also be completed using XenCenter by editing the VM configuration and connecting the existing VIFs of the VM to the bond network.
11. Repeat steps 3 - 10 for other hosts.
12. Restart the VMs previously shut down.

Configuring a dedicated storage NIC

XenServer allows use of either XenCenter or the xe CLI to configure and dedicate a NIC to specific functions, such as storage traffic.

Assigning a NIC to a specific function will prevent the use of the NIC for other functions such as host management, but requires that the appropriate network configuration be in place in order to ensure the NIC is used for the desired traffic. For example, to dedicate a NIC to storage traffic the NIC, storage target, switch, and/or VLAN must be configured such that the target is only accessible over the assigned NIC. This allows use of standard IP routing to control how traffic is routed between multiple NICs within a XenServer.

Note:

Before dedicating a network interface as a storage interface for use with iSCSI or NFS SRs, ensure that the dedicated interface uses a separate IP subnet which is not routable from the main management interface. If this is not enforced, then storage traffic may be directed over the main management interface after a host reboot, due to the order in which network interfaces are initialized.

To assign NIC functions using the xe CLI

1. Ensure that the PIF is on a separate subnet, or routing is configured to suit your network topology in order to force the desired traffic over the selected PIF.
2. Setup an IP configuration for the PIF, adding appropriate values for the mode parameter and if using static IP addressing the IP, netmask, gateway, and DNS parameters:

```
xe pif-reconfigure-ip mode=<DHCP | Static> uuid=<pif-uuid>
```

3. Set the PIF's disallow-unplug parameter to true:

```
xe pif-param-set disallow-unplug=true uuid=<pif-uuid>
```

```
xe pif-param-set other-config:management_purpose="Storage" uuid=<pif-uuid>
```

If you want to use a storage interface that can be routed from the management interface also (bearing in mind that this configuration is not recommended), then you have two options:

- After a host reboot, ensure that the storage interface is correctly configured, and use the **xe pbd-unplug** and **xe pbd-plug** commands to reinitialize the storage connections on the host. This will restart the storage connection and route it over the correct interface.
- Alternatively, you can use **xe pif-forget** to remove the interface from the XenServer database, and manually configure it in the control domain. This is an advanced option and requires you to be familiar with how to manually configure Linux networking.

Controlling Quality of Service (QoS)

Citrix XenServer allows an optional Quality of Service (QoS) value to be set on VM virtual network interfaces (VIFs) using the CLI. The supported QoS algorithm type is rate limiting, specified as a maximum transfer rate for the VIF in Kb per second.

For example, to limit a VIF to a maximum transfer rate of 100kb/s, use the **vif-param-set** command:

```
xe vif-param-set uuid=<vif-uuid> qos_algorithm_type=ratelimit
xe vif-param-set uuid=<vif-uuid> qos_algorithm_params:kpbs=100
```

Changing networking configuration options

This section discusses how to change the networking configuration of a XenServer host. This includes:



- changing the hostname
- adding or removing DNS servers
- changing IP addresses
- changing which NIC is used as the management interface
- adding a new physical NIC to the server

Hostname

The system hostname is defined in the pool-wide database and modified using the **xe host-set-hostname-live** CLI command as follows:

```
xe host-set-hostname-live uuid=<host_uuid> host-name=example
```

The underlying control domain hostname changes dynamically to reflect the new hostname.

DNS servers

To add or remove DNS servers in the IP addressing configuration of a XenServer host, use the **pif-reconfigure-ip** command. For example, for a PIF with a static IP:

```
pif-reconfigure-ip uuid=<pif_uuid> mode=static DNS=<new_dns_ip>
```

Changing IP address configuration for a standalone host

Network interface configuration can be changed using the xe CLI. The underlying network configuration scripts should not be modified directly.

To modify the IP address configuration of a PIF, use the **pif-reconfigure-ip** CLI command. See [the section called “pif-reconfigure-ip”](#) for details on the parameters of the **pif-reconfigure-ip** command.

Note:

See [the section called “Changing IP address configuration in resource pools”](#) for details on changing host IP addresses in resource pools.

Changing IP address configuration in resource pools

XenServer hosts in resource pools have a single management IP address used for management and communication to and from other hosts in the pool. The steps required to change the IP address of a host's management interface are different for master and other hosts.

Note:

Caution should be used when changing the IP address of a server, and other networking parameters. Depending upon the network topology and the change being made, connections to network storage may be lost. If this happens the storage must be replugged using the **Repair Storage** function in XenCenter, or the **pbid-plug** command using the CLI. For this reason, it may be advisable to migrate VMs away from the server before changing its IP configuration.

Changing the IP address of a pool member host

1. Use the **pif-reconfigure-ip** CLI command to set the IP address as desired. See [Appendix A, Command line interface](#) for details on the parameters of the **pif-reconfigure-ip** command:

```
xe pif-reconfigure-ip uuid=<pif_uuid> mode=DHCP
```

2. Use the **host-list** CLI command to confirm that the member host has successfully reconnected to the master host by checking that all the other XenServer hosts in the pool are visible:

```
xe host-list
```

Changing the IP address of the master XenServer host requires additional steps because each of the member hosts uses the advertised IP address of the pool master for communication and will not know how to contact the master when its IP address changes.

Whenever possible, use a dedicated IP address that is not likely to change for the lifetime of the pool for pool masters.

To change the IP address of a pool master host

1. Use the **pif-reconfigure-ip** CLI command to set the IP address as desired. See [Appendix A, Command line interface](#) for details on the parameters of the **pif-reconfigure-ip** command:

```
xe pif-reconfigure-ip uuid=<pif_uuid> mode=DHCP
```

2. When the IP address of the pool master host is changed, all member hosts will enter into an emergency mode when they fail to contact the master host.
3. On the master XenServer host, use the **pool-recover-slaves** command to force the master to contact each of the member hosts and inform them of the new master IP address:

```
xe pool-recover-slaves
```

Management interface

When XenServer is installed on a host with multiple NICs, one NIC is selected for use as the management interface. The management interface is used for XenCenter connections to the host and for host-to-host communication.

To change the NIC used for the management interface

1. Use the **pif-list** command to determine which PIF corresponds to the NIC to be used as the management interface. The UUID of each PIF is returned.

```
xe pif-list
```

2. Use the **pif-param-list** command to verify the IP addressing configuration for the PIF that will be used for the management interface. If necessary, use the **pif-reconfigure-ip** command to configure IP addressing for the PIF to be used. See [Appendix A, Command line interface](#) for more detail on the options available for the **pif-reconfigure-ip** command.

```
xe pif-param-list uuid=<pif_uuid>
```

3. Use the **host-management-reconfigure** CLI command to change the PIF used for the management interface. If this host is part of a resource pool, *this command must be issued on the member host console*:

```
xe host-management-reconfigure pif-uuid=<pif_uuid>
```

Warning:

Putting the management interface on a VLAN network is not supported.

Disabling management access

To disable remote access to the management console entirely, use the **host-management-disable** CLI command.

Warning:

Once the management interface is disabled, you will have to log in on the physical host console to perform management tasks and external interfaces such as XenCenter will no longer work.

Adding a new physical NIC

Install a new physical NIC on a XenServer host in the usual manner. Then, after restarting the server, run the `xe CLI` command **pif-scan** to cause a new PIF object to be created for the new NIC.

NIC/PIF ordering in resource pools

It is possible for physical NIC devices to be discovered in different orders on different servers even though the servers contain the same hardware. Verifying NIC ordering is recommended before using the pooling features of XenServer.

Verifying NIC ordering

Use the **pif-list** command to verify that NIC ordering is consistent across your XenServer hosts. Review the MAC address and carrier (link state) parameters associated with each PIF to verify that the devices discovered (`eth0`, `eth1`, etc.) correspond to the appropriate physical port on the server.

```
xe pif-list params=uuid,device,MAC,currently-attached,carrier,management, \
IP-configuration-mode
```

```
uuid ( RO)                : 1ef8209d-5db5-cf69-3fe6-0e8d24f8f518
      device ( RO): eth0
      MAC ( RO): 00:19:bb:2d:7e:8a
      currently-attached ( RO): true
      management ( RO): true
      IP-configuration-mode ( RO): DHCP
      carrier ( RO): true
```

```
uuid ( RO)                : 829fd476-2bbb-67bb-139f-d607c09e9110
      device ( RO): eth1
      MAC ( RO): 00:19:bb:2d:7e:7a
      currently-attached ( RO): false
      management ( RO): false
      IP-configuration-mode ( RO): None
      carrier ( RO): true
```

If the hosts have already been joined in a pool, add the `host-uuid` parameter to the **pif-list** command to scope the results to the PIFs on a given host.

Re-ordering NICs

It is not possible to directly rename a PIF, although you can use the **pif-forget** and **pif-introduce** commands to achieve the same effect with the following restrictions:

- The XenServer host must be standalone and not joined to a resource pool.
- Re-ordering a PIF configured as the management interface of the host requires additional steps which are included in the example below. Because the management interface must first be disabled the commands must be entered directly on the host console.

For the example configuration shown above use the following steps to change the NIC ordering so that `eth0` corresponds to the device with a MAC address of `00:19:bb:2d:7e:7a`:

1. Use XenCenter or the **vm-shutdown** command to shut down all VMs in the pool to force existing VIFs to be unplugged from their networks.

```
xe vm-shutdown uuid=<vm_uuid>
```

2. Use the **host-management-disable** command to disable the management interface:

```
xe host-management-disable
```

3. Use the **pif-forget** command to remove the two incorrect PIF records:

```
xe pif-forget uuid=1ef8209d-5db5-cf69-3fe6-0e8d24f8f518
xe pif-forget uuid=829fd476-2bbb-67bb-139f-d607c09e9110
```

4. Use the **pif-introduce** command to re-introduce the devices with the desired naming:

```
xe pif-introduce device=eth0 host-uuid=<host_uuid> mac=00:19:bb:2d:7e:7a
xe pif-introduce device=eth1 host-uuid=<host_uuid> mac=00:19:bb:2d:7e:8a
```

5. Use the **pif-list** command again to verify the new configuration:

```
xe pif-list params=uuid,device,MAC
```

6. Use the **pif-reconfigure-ip** command to reset the management interface IP addressing configuration. See [Appendix A, Command line interface](#) for details on the parameters of the **pif-reconfigure-ip** command.

```
xe pif-reconfigure-ip uuid=<728d9e7f-62ed-a477-2c71-3974d75972eb> mode=dhcp
```

7. Use the **host-management-reconfigure** command to set the management interface to the desired PIF and re-enable external management connectivity to the host:

```
xe host-management-reconfigure pif-uuid=<728d9e7f-62ed-a477-2c71-3974d75972eb>
```

Networking Troubleshooting

If you are having problems with configuring networking, first ensure that you have not directly modified any of the control domain `ifcfg-*` files directly. These files are directly managed by the control domain host agent, and changes will be overwritten.

Diagnosing network corruption

Some network card models require firmware upgrades from the vendor to work reliably under load, or when certain optimizations are turned on. If you are seeing corrupted traffic to VMs, then you should first try to obtain the latest recommended firmware from your vendor and apply a BIOS update.

If the problem still persists, then you can use the CLI to disable receive / transmit offload optimizations on the physical interface.

Warning:

Disabling receive / transmit offload optimizations can result in a performance loss and / or increased CPU usage.

First, determine the UUID of the physical interface. You can filter on the *device* field as follows:

```
xe pif-list device=eth0
```

Next, set the following parameter on the PIF to disable TX offload:

```
xe pif-param-set uuid=<pif_uuid> other-config:ethtool-tx=off
```



Finally, re-plug the PIF or reboot the host for the change to take effect.

Recovering from a bad network configuration

In some cases it is possible to render networking unusable by creating an incorrect configuration. This is particularly true when attempting to make network configuration changes on a member XenServer host.

If a loss of networking occurs, the following notes may be useful in recovering and regaining network connectivity:

- Citrix recommends that you ensure networking configuration is set up correctly before creating a resource pool, as it is usually easier to recover from a bad configuration in a non-pooled state.
- The **host-management-reconfigure** and **host-management-disable** commands affect the XenServer host on which they are run and so are not suitable for use on one host in a pool to change the configuration of another. Run these commands directly on the console of the XenServer host to be affected, or use the **xe** *-s*, *-u*, and *-pw* remote connection options.
- When the `xapi` service starts, it will apply configuration to the management interface first. The name of the management interface is saved in the `/etc/xensource-inventory` file. In extreme cases, you can stop the `xapi` service by running **service xapi stop** at the console, edit the inventory file to set the management interface to a safe default, and then ensure that the `ifcfg` files in `/etc/sysconfig/network-scripts` have correct configurations for a minimal network configuration (including one interface and one bridge; for example, `eth0` on the `xenbr0` bridge).

Workload Balancing

What's New?

This topic describes new features and changes from the previous release of Workload Balancing.

New Features

Workload Balancing now provides a series of features that work together to facilitate reducing power during off-peak work hours. This includes functionality that automatically adjusts your Workload Balancing placement strategy during off-peak periods, automatically rebalances virtual machines to compress workloads onto hosts as densely as possible, and automatically turns off power on those lightly loaded machines after their workloads have been relocated.

This functionality is provided through the following features:

- **Power Management.** Power Management lets you configure Workload Balancing to consolidate workloads on the lowest possible number of physical servers (hosts) and turn the power off on unused hosts automatically when their capacity is not required. It can also turn the power back on for hosts that it has previously turned off when their capacity is needed. This feature can be useful, for example, if you want to power off unused servers at night to save energy. See also [the section called “Optimizing and Managing Power Automatically”](#).
- **Automation.** You can configure Workload Balancing to apply optimization recommendations automatically so that you do not need to accept recommendations manually. See also [the section called “Optimizing and Managing Power Automatically”](#).
- **Schedule Different Optimization Modes for Different Times of Day.** You can now schedule optimization modes to switch automatically between Maximum Density to Maximum Performance for the time periods you specify. Therefore, if you want Workload Balancing to optimize workloads for Maximum Performance during your organization busiest periods (for example, daytime) and Maximum Density at night, you can configure it to do so automatically using a schedule that can vary by the day of the week.
- **Host Exclusion.** The ability to exclude a specific host in a pool from Workload Balancing's optimization recommendations even though Workload Balancing is enabled for all other hosts in the pool. See also [the section called “Optimizing and Managing Power Automatically”](#).

Other new features include:

- **Audit Log Report.** You can generate reports for changes specific user accounts, including the Local Super User (LSU), made to a pool for auditing purposes.
- **Enhanced Reporting.** This release includes an optimized user interface for reporting as well as additional reports, including the Audit Log report.
- **Report Subscriptions.** You can now have a report generated and delivered to your email automatically at preset dates and times.
- **Usability Enhancements.**
 - **Simplified Installation.** You no longer need to install the Microsoft SQL Server 2005 Backward Compatibility Components before installation.
 - **Improved User Interface.** Features such as Initialization and Configuration are streamlined and easier to use.



- **Improved Error Messages.** Improved error messages make issues easier to understand. Also, error messages are now documented with suggested actions.

In addition, this release of Workload Balancing is aware of the XenServer Dynamic Memory Control feature — Workload Balancing considers a virtual machine's memory range and target when providing virtual-machine placement, host evacuation, and pool optimization recommendations.

To see the new features included in this release of Workload Balancing, you must also upgrade XenCenter and XenServer to the latest versions.

Changes

This release of Workload Balancing also contains the following changes:

- **Installation Changes.**
 - **Prerequisites.**
 - Microsoft SQL Server 2005 Backward Compatibility Components are no longer required.
 - Changes to User Interface.
 - Protocol Selection Removed. You no longer need to select the protocol; it is selected for you silently during installation. HTTPS is and remains the only supported protocol.
 - Certificate configuration is no longer required. However, you can configure Trusted Certificates for Workload Balancing to use with XenServer after installation using the new configuration tool.
- **Full Initial Configuration No Longer Required to Run.** You no longer need to configure optimization and placement settings when you initialize Workload Balancing. Instead, if desired, you modify default placement settings after initialization in a separate properties dialog.
- **New Dialogs for Configuration and Initialization.** Server settings (initialization) and configuration settings are now in two separate dialog boxes: the **Connect to WLB Server** dialog box and the **Workload Balancing Configuration** properties dialog replace the **Configure Workload Balancing** wizard.
- **Overcommit CPU.** The Overcommit CPU feature is not available in this Workload Balancing release. If you perform an upgrade using the installation wizard and you enabled the feature, the installation wizard will automatically disable it for you.

Workload Balancing Overview

Workload Balancing is a XenServer feature that balances virtual machines and locates them on the best possible servers for their workload in a resource pool. For example, Workload Balancing can:

- Balance virtual-machine workloads across hosts in a XenServer resource pool
- Determine the best host on which to start a virtual machine
- Determine the best host on which to resume a virtual machine that you powered off
- Determine the best host to move a virtual machine to when a host fails
- Determine the optimal server for each of the host's virtual machines when you put a host into or take a host out of Maintenance mode

Depending on your preference, Workload Balancing can accomplish these tasks automatically or prompt you to accept its rebalancing and placement recommendations. You can also configure Workload Balancing to power off hosts automatically at specific times of day (for example, to save power at night).



Workload Balancing works by evaluating the utilization of workloads across a pool and, when a host exceeds its threshold, relocating it to a less-taxed host in the pool. To rebalance workloads, Workload Balancing moves workloads to keep hosts in balance.

To ensure the rebalancing and placement recommendations align with your environment's needs, you can configure Workload Balancing to optimize workloads for either resource performance or to maximize density. These optimization modes can be configured to change automatically at predefined times or stay the same at all times. For additional granularity, you can fine-tune the weighting of individual resource metrics (CPU, network, memory, and disk).

To help you perform capacity planning, Workload Balancing provides historical reports about host and pool health, optimization and virtual-machine performance, and virtual-machine motion history.

Workload Balancing Basic Concepts

Workload Balancing captures data for resource performance on virtual machines and physical hosts. It uses this data, combined with the preferences you set, to provide optimization and placement recommendations. Workload Balancing stores performance data in a SQL Server database: the longer Workload Balancing runs, the more precise its recommendations become.

Workload Balancing recommends moving virtual-machine workloads across a pool to get the maximum efficiency, which means either *performance* or *density* depending on your goals. Within a Workload Balancing context:

- **Performance** refers to the usage of physical resources on a host (for example, the CPU, memory, network, and disk utilization on a host). When you set Workload Balancing to maximize performance, it recommends placing virtual machines to ensure the maximum amount of resources are available for each virtual machine.
- **Density** refers to the number of virtual machines on a host. When you set Workload Balancing to maximize density, it recommends placing virtual machines to ensure they have adequate computing power so you can reduce the number of hosts powered on in a pool.

Workload Balancing does not conflict with settings you already specified for High Availability. Citrix designed the features to work in conjunction with each other.

Workload Balancing Installation Overview

Workload Balancing is typically deployed on the same computer as the data store or with the data store on a separate database server.

Installing Workload Balancing requires that you:

1. Install SQL Server or SQL Server Express.
2. Install Workload Balancing (on physical computer or a virtual machine).

Typically, you install and configure Workload Balancing after you have created one or more XenServer resource pools in your environment.

You can install Workload Balancing in one of two ways:

- **Installation Wizard.** Start the installation wizard from Setup.exe. Citrix suggests installing Workload Balancing from the installation wizard because this method checks your system meets the installation requirements.



- **Command Line.** If you install Workload Balancing from the command line, the prerequisites are not checked. For Msiexec properties, see [CTX124480](#).

When you install Workload Balancing, Setup creates the data store by installing it across the network onto the database. (You do not need to run Workload Balancing Setup locally on the database server.)

After installation, you must set up Workload Balancing server communication settings before you can use Workload Balancing to optimize workloads. For information, see [the section called “Initializing Workload Balancing”](#).

Note:

To install Workload Balancing on multiple servers, see [CTX124289](#).

For information about System Requirements, see [the section called “Workload Balancing System Requirements”](#). For installation instructions, see [the section called “Installing Workload Balancing”](#).

Workload Balancing System Requirements

This topic lists the prerequisites for installing Workload Balancing or upgrading it to this release.

For information about data store requirements, see [the section called “Workload Balancing Data Store Requirements”](#).

Important:

Citrix recommends installing the Workload Balancing server on Windows Server 2008 R2 since it simplifies downloading and installing the recommended prerequisites.

Supported XenServer Versions

- XenServer 5.6

Supported Operating Systems

Unless otherwise noted, Workload Balancing components run on the following operating systems (32-bit and 64-bit editions):

- Windows Server 2008 R2
- Windows Server 2008
- Windows Server 2003, Service Pack 2
- Windows 7
- Windows Vista
- Windows XP Professional, Service Pack 2 or Service Pack 3

Recommended Hardware

Unless otherwise noted, Workload Balancing requires the following hardware:

- CPU: 2GHz or faster



- Memory: 2GB recommended (1GB of RAM required)
- Disk Space:
 - SQL Server Express: 4GB
 - SQL Server: 20GB (approximate amount consumed)

Citrix recommends that the Workload Balancing server have a minimum of a dual-core processor.

Workload Balancing Data Store Requirements

This topic provides information about the SQL Server versions and configurations that Workload Balancing supports. It also provides information about additional compatibility and authentication requirements.

In addition to the prerequisites SQL Server and SQL Server Express require, the data store requires the following:

Note:

In this topic, the term *SQL Server* refers to both SQL Server and SQL Server Express unless the version is mentioned explicitly.

Database

The 32-bit or 64-bit edition of:

- SQL Server 2008 Express. The 32-bit edition of SQL Server 2008 Express, SP1 is available on the Workload Balancing installation media in the **sql** folder. To install it, run `en_sql_server_2008_express_with_tools_x86`.
- SQL Server 2008 (Standard edition or better)
- SQL Server 2005, Service Pack 1 or higher (Standard edition or better)

Note:

Windows Server 2008 servers require SQL Server 2005, Service Pack 2 or higher.

Required Configurations

- Configure SQL Server for case-insensitive collation. Workload Balancing does not currently support case-sensitive collation.

Hard Drive Space

- SQL Server Express: 4GB
- SQL Server: 20GB

SQL Server Database Authentication Requirements

During installation, Setup must connect and authenticate to the database server to create the data store. Configure the SQL Server database instance to use either:

- **Windows Authentication mode**, or



- **SQL Server and Windows Authentication mode** (Mixed Mode authentication)

If you create an account on the database for use during Setup, the account must have **sysadmin** privileges for the database instance where you want to create the Workload Balancing data store.

Operating System Language Support

Workload Balancing is supported on the following operating system languages:

- US English
- Japanese (Native JP)

Note:

When Workload Balancing is installed in a distributed, multi-server deployment, as described in [CTX124289](#), the operating system languages must match on both computers if the database and Web servers are installed on separate computers.

Preinstallation Considerations

You may need to configure software in your environment so that Workload Balancing can function correctly. Review the following considerations and determine if they apply to your environment. Also, check the XenServer Release Notes for additional, late-breaking, release-specific requirements.

- **Internet Connectivity.** Workload Balancing requires Internet connectivity for digital signature verification. If you do not have Internet connectivity, see [XenServer Workload Balancing Installation Fails Without Available Internet Connection], [CTX123074](#).
- **Account for Workload Balancing.** During Setup, you must specify the authorization type (a single user or group) and enter the user or group with permissions to make requests of the Web Service Host service. For additional information, see [the section called "Authorization for Workload Balancing"](#). The user can be either a domain account or an account local to the computer running Workload Balancing (or the Web Service Host service). Likewise, groups can also be domain or local.

Before Setup, you may want to create a user account or user group for XenServer to use to connect to Workload Balancing.

Important:

When you create this account in Windows, Citrix suggests enabling the **Password never expires** option.

- **SSL/TLS Certificate.** XenServer and Workload Balancing communicate over HTTPS. Consequently, during Workload Balancing Setup, Workload Balancing automatically creates a self-signed certificate on your behalf.
- **Group Policy.** If the server on which you are installing Workload Balancing is a member of a Group Policy Organizational Unit, ensure that current or scheduled, future policies do not prohibit Workload Balancing or its services from running.
- **Communications Port.** Determine the port over which you want the Workload Balancing server to communicate. You will be prompted for this port during Setup. By default, Workload Balancing server uses 8012.
- **Synchronizing Time on Host and VM.** When installing Workload Balancing on a VM, it is essential that the time on the physical server hosting the WLB VM and the WLB VM itself match. You might want to consider pointing both the host and VM to a Network Time (NTP) server.



WLB Access Control Permissions

When Role Based Access Control (RBAC) is implemented in your environment, all user roles can display the WLB tab. However, not all roles can perform all operations. The following table lists the *minimum* role administrators require to use Workload Balancing features:

Task	Minimum Required Role
Configure, Initialize, Enable, Disable WLB	Pool Operator
Apply WLB Optimization Recommendations (in WLB tab)	Pool Operator
Modify WLB Report Subscriptions	Pool Operator
Accept WLB Placement Recommendations ("star" recommendations)	VM Power Admin
Generate WLB Reports, including the Pool Audit Trail report	VM Operator
Display WLB Configuration	Read Only

If an administrator tries to use Workload Balancing and he or she does not have sufficient permissions, a role elevation dialog appears. For more information about RBAC, see [the section called "Role Based Access Control"](#).

Note:

Generating WLB reports requires the VM Operator role or higher. Users assigned the Read Only role cannot generate reports in this release.

Installing Workload Balancing

Before installing Workload Balancing, you must:

1. Install a SQL Server or SQL Server Express database as described in [the section called "Workload Balancing Data Store Requirements"](#).
2. Have a login on the SQL Server database instance that has SQL Login creation privileges. For SQL Server Authentication, the account needs **sysadmin** privileges.
3. Create an account for Workload Balancing, as described in [the section called "Preinstallation Considerations"](#) and have its name on hand.
4. Configure all Workload Balancing servers to meet the system requirements described in [the section called "Workload Balancing System Requirements"](#).

After Setup is finished installing Workload Balancing, verify that it installed correctly. See [the section called "To verify your Workload Balancing installation"](#). Then you must set up the connection between Workload Balancing and your pool. See [the section called "Initializing Workload Balancing"](#) before Workload Balancing can begin gathering data and making recommendations.

Note:

Workload Balancing requires Internet connectivity for digital signature verification as described in [the section called "Preinstallation Considerations"](#).



To install Workload Balancing server

The following procedure installs Workload Balancing on one computer:

1. Launch the **Workload Balancing Setup** wizard from Autorun.exe, and select the **Workload Balancing** option.

Note:

If you are installing on Windows Vista, Windows Server 2008, or Windows 7 and the User Account Control (UAC) is enabled, right-click Autorun.exe and select **Run as administrator**. (To install Workload Balancing with UAC enabled, you must have sufficient administrative privileges to authorize an elevation and install software.)

2. After the initial Welcome page appears, click **Next**.
3. In the **Setup Type** page, select **Workload Balancing Services and Data Store**, and click **Next**. This option lets you install Workload Balancing and configure its data store. After you click Next, Workload Balancing Setup verifies that your system has the correct prerequisites.
4. Accept the End-User License Agreement.
5. In the **Component Selection** page, select **Database** and all of the **Services**. This is the default selection.

Note:

If you are reinstalling Workload Balancing and you have an existing remote database from this release, deselect the database option in the **Component Selection** page of Setup. If you do not clear this option, then Setup attempts to reinstall a database to the location you specify in the **Database Server** page and issues an error.

6. In the **Database Server** page, in the **SQL Server Selection** section, select one of the following:
 - **Enter the name of a database server**. Lets you type the name of the database server that will host the data store. Specify an instance name by appending it to the database name (for example, **database-name\yourinstancename**).

Note:

By default, SQL Express creates an instance name even if you do not specify one during SQL Express installation. If you do not explicitly specify an instance name, SQL Express Setup creates an instance on the database named "sql express."

If you installed SQL Express and specified an instance name, append the server name with **\yourinstancename**. If you installed SQL Express without specifying an instance name, append the server name with **\sql express**.

- **Choose an existing database server**. Lets you select the database server from a list of servers Workload Balancing Setup detected on your network.

Note:

Workload Balancing installation may fail if you specify a name for the data store that uses certain Japanese characters.

7. In the **Install Using** section, select one of the following methods of authentication:
 - **Windows Authentication**. This option uses your current credentials (that is, the Windows credentials you used to log on to the computer on which you are installing Workload Balancing). To select this option, your current Windows credentials must have been added as a login to the SQL Server database server (instance).



- **SQL Server Authentication.** To select this option, you must have configured SQL Server to support Mixed Mode authentication.

Note:

Citrix recommends clicking **Test Connect** to ensure Setup can use the credentials you provided to contact the database server.

8. In the **Database Information** page, select **Install a new Workload Balancing data store** and type the name you want to assign to the Workload Balancing database in SQL Server. The default database name is `WorkloadBalancing`.
9. In the **Web Service Host Account Information** page, edit the port number, if necessary. The port is set to 8012 by default.

Note:

If you change the port here, you must also change it on XenServer using either the **Connect to WLB Server** dialog or the XE commands. Make sure that the port you specify for the Web Service is open in any firewalls, including the Windows Firewall.

10. For the account (on the Workload Balancing server) that XenServer will use to connect to Workload Balancing, select the authorization type, **User** or **Group**, and type one of the following:

- **User name.** Enter the name of the account you created for XenServer (for example, `workloadbalancing_user`).
- **Group name.** Enter the group name containing the account you want XenServer to use.

Specifying the authorization type lets Workload Balancing recognize the XenServer's connection. For more information, see [the section called "Authorization for Workload Balancing"](#). You do not specify the password until you configure Workload Balancing.

Note:

Specifying the Domain Users Group as the authorization type (that is, the group) in Workload Balancing installation is not supported. If you use this group, Workload Balancing cannot be successfully initialized.

11. Click **Install**.

To verify your Workload Balancing installation

Workload Balancing Setup does not install an icon in the Windows Start menu. Use this procedure to verify that Workload Balancing installed correctly before trying to connect to the Workload Balancing server using the **Connect to WLB Server** dialog.

1. Verify Windows Add or Remove Programs lists Citrix Workload Balancing as a currently installed programs.
2. Check for the following services in the Windows Services panel:
 - Citrix WLB Analysis Engine
 - Citrix WLB Data Collection Manager
 - Citrix WLB Web Service Host

All of these services must be started and running before you start configuring Workload Balancing.

3. If Workload Balancing services are not in the Services panel, check the installation log to see if it installed successfully:



- If you used the Setup wizard, the log is at C:\Documents and Settings\username\Local Settings\Temp\msibootstrapper2CSM_MSI_Install.log (by default). On Windows Vista and Windows Server 2008, this log is at C:\Users\username\AppData\Local\Temp\msibootstrapper2CSM_MSI_Install.log. User name is the name of the user logged on during installation.
- If you used the Setup properties (Msiexec), the log is at C:\log.txt (by default) or wherever you specified for Setup to create it.

Configuring Firewalls

If Workload Balancing server is behind a firewall, you must configure the **WLB Web Service HTTPS Port** rule in the Windows Firewall settings (specifically, the Private Windows firewall profile) for a Domain and Public profile.

By default Workload Balancing Setup creates a Firewall rule ("WLB Web Service HTTPS Port") in the "Private" Windows Firewall Profile. However, inbound communication via the "Domain" and "Public" profile may also be necessary for WLB to function properly. To configure this setting, in the existing WLB Firewall rule, also enable the **Domain** and **Public** profile check boxes.

Note:

The following error appears if the Workload Balancing server is behind a firewall and you did not configure the appropriate firewall settings: "There was an error connecting to the Workload Balancing server: <pool name> Click Initialize WLB to reinitialize the connection settings." This error may also appear if the Workload Balancing server is otherwise unreachable.

Upgrading Workload Balancing

To see the new features included in this release of Workload Balancing, you must also upgrade XenCenter and XenServer to the latest versions.

- If you do not upgrade *XenCenter*, the new features will not appear in the XenCenter user interface. However, you will be able to use them through the XE commands.
- If you do not upgrade *XenServer*, you will not be able to use new Workload Balancing features; however, upgrading the Workload Balancing server will install fixes for known issues.

Note:

Citrix recommends backing up your Workload Balancing database before upgrading the Workload Balancing server. Optionally, you may want to restart your Workload Balancing server so that all Workload Balancing processes are restarted, providing a known state from which to upgrade. Also, upgrade your XenServer hosts before upgrading the Workload Balancing server.

You can upgrade directly to the latest release of Workload Balancing (Workload Balancing 2.0) from the previous version (Workload Balancing 1.0).

Before upgrading, Citrix recommends backing up your data using SQL Management Studio. When you upgrade, you are prompted for the location of the existing Workload Balancing data store. After you specify it, the installation wizard upgrades it to the newest data store automatically.

The MSI file checks for existing versions of the Workload Balancing database and Workload Balancing components each time you launch the installation wizard. Upgrades from the XenServer 5.6 Alpha and Beta releases are not supported.

Note:



The time it takes to upgrade the server may be affected by the distribution of the database file group. Additional time may be required to perform file group moves during the upgrade. In addition, the size of the database affects the time required to perform an upgrade.

Upgrading Workload Balancing on the Same Operating System

If you only want to upgrade the Workload Balancing server software and you are not upgrading operating system on its server, do the following:

1. Temporarily disable Workload Balancing, as described in [the section called “Disabling Workload Balancing”](#).
2. Install any upgrades to system requirements that are new in this release.
3. Restart the VM or server where Workload Balancing is installed.
4. Run the Workload Balancing Installation Wizard. The Workload Balancing Installation Wizard automatically detects the presence of an existing installation of Workload Balancing and prompts you to upgrade that installation.

During upgrade, the installation wizard displays a subset of the pages it displays during a new installation. For information about the pages in the installation wizard, see [the section called “To install Workload Balancing server”](#).

Upgrading SQL Server

If you want to upgrade the SQL Server version (for example, move from SQL Server Express 2005 to SQL Server Express 2008), you can use the SQL Management Studio tools to migrate your data.

Upgrading Workload Balancing and the Operating System

If you want to migrate to a newer version of Windows (for example, from Windows Server 2003 to Windows Server 2008 R2), you cannot upgrade Workload Balancing using the Upgrade mode in the installation wizard. Instead, you must reimage the server with the new operating system and then install Workload Balancing.

If your data store is on the same server as Workload Balancing, use SQL Management Studio tools to save your data to another location before reimaging the server and installing the new operating system.

Initializing Workload Balancing

Following Workload Balancing Setup, you must enable (that is, initialize) Workload Balancing on each resource pool you want to monitor before Workload Balancing can gather data for that pool. Initializing Workload Balancing requires that you:

- **Specify the Workload Balancing server** you want the resource pool to use and its port number.
- **Specify the credentials** for communications, including the credentials XenServer will use to connect to the Workload Balancing server and Workload Balancing will use to connect to XenServer, as described in [the section called “Authorization for Workload Balancing”](#).

When you first enable Workload Balancing, it uses the default configuration settings for optimizing workload. Automatic features, such as Automated Optimization Mode, Power Management, and Automation, are disabled by default.

Important:



After Workload Balancing is running for a period of time, if you do not receive optimal placement recommendations, Citrix strongly recommends you evaluate your performance thresholds as described in [the section called “Using Workload Balancing Reports for Tasks”](#). It is critical to set Workload Balancing to the correct thresholds for your environment or its recommendations might not be appropriate.

Before initializing Workload Balancing, configure your antivirus software to exclude Workload Balancing folders, as described in [the section called “Configuring Antivirus Software”](#).

Note:

After the initial configuration, the **Initialize** button on the **WLB** tab changes to a **Disable** button. This is because after initialization you cannot modify the Workload Balancing server a resource pool uses without disabling Workload Balancing on that pool and then reconfiguring it. For information, see [the section called “Reconfiguring a Pool to Use Another WLB Server”](#).

You can use the **Connect to WLB Server** dialog box in XenCenter or the XE commands to initialize Workload Balancing.

To initialize Workload Balancing

Use this procedure to enable and perform the initial configuration of Workload Balancing for a resource pool. To complete this dialog, you need the:

- IP address (or NetBIOS name) and port of the Workload Balancing server
- Credentials for the resource pool you want Workload Balancing to monitor
- Credentials for the account you created on the Workload Balancing server

Note:

Before the Workload Balancing feature can begin collecting performance data, the XenServers you want to balance must be part of a resource pool.

1. In the **Resources** pane of XenCenter, select `XenCenter > your-resource-pool`.
2. In the **Properties** pane, click the **WLB** tab.
3. In the **WLB** tab, click **Initialize WLB**.

The **Connect to WLB Server** dialog box appears.

4. In the **Server Address** section, dialog box, enter the following:
 - a. In the **Address** box, type the IP address or NetBIOS name of the Workload Balancing server. You can also enter a fully qualified domain name (FQDN).
 - b. (Optional.) Edit the port number if you want XenServer to connect to Workload Balancing using a different port. Entering a new port number here sets a different communications port on the Workload Balancing server.

By default, XenServer connects to Workload Balancing (specifically the Web Service Host service) on port 8012.

Note:

Do not edit this port number unless you have changed it during Workload Balancing Setup. The port number value specified during Setup and in the **Workload Balancing Configuration** dialog must match.

5. In the **WLB Server Credentials** section, enter the user name (for example, `workloadbalancing_user`) and password the computers running XenServer will use to connect to the Workload Balancing server.

This must be the account or group that was configured during the installation of the Workload Balancing Server. For information, see [the section called “Authorization for Workload Balancing”](#).

6. In the **XenServer Credentials** section, enter the user name and password for the pool you are configuring (typically the password for the pool master). Workload Balancing will use these credentials to connect to the computers running XenServer in that pool.

To use the credentials with which you are currently logged into XenServer, select the **Use the current XenCenter credentials** check box. If you have assigned a role to this account using the Access Control feature (RBAC), be sure the role has sufficient permissions to configure Workload Balancing. See [the section called “WLB Access Control Permissions”](#).

7. After initialization, if you want to modify advanced settings for thresholds and change the priority given to specific resources, see [the section called “Configuring Workload Balancing Settings”](#).

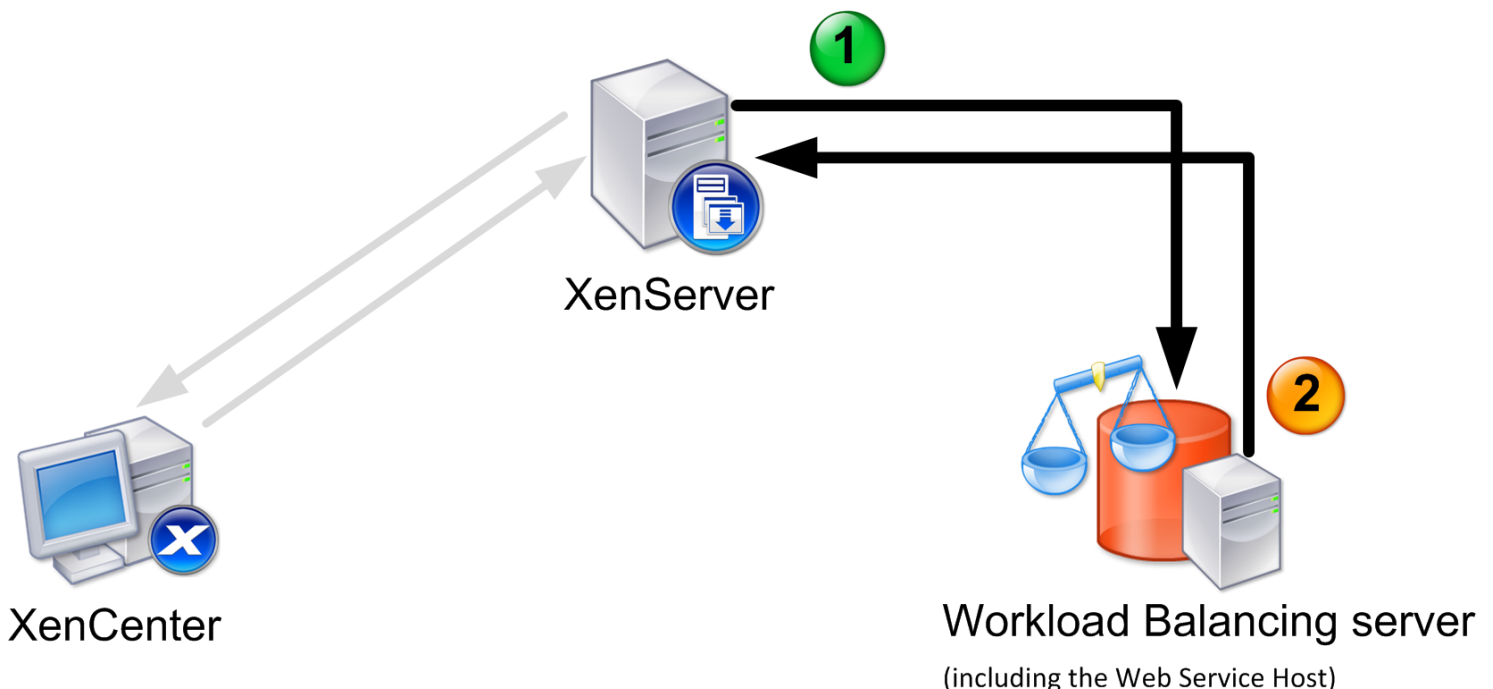
Authorization for Workload Balancing

When you are configuring a XenServer resource pool to use Workload Balancing, you must specify credentials for two accounts:

- **User Account for Workload Balancing to connect to XenServer.** Workload Balancing uses a XenServer user account to connect to XenServer. You provide Workload Balancing with this account's credentials when you run the **Connect to WLB Server** dialog. Typically, you specify the credentials for the pool (that is, the pool master's credentials).
- **User Account for XenServer to Connect to Workload Balancing.** XenServer communicates with the Web Service Host using the user account you created before Setup.

During Workload Balancing Setup, you specified the authorization type (a single user or group) and the user or group with permissions to make requests from the Web Service Host service.

During initialization, you must provide XenServer with this account's credentials when you configure the connection settings in the **Connect to WLB Server** dialog.





Configuring Antivirus Software

By default, most antivirus programs are configured to scan all files on the hard disk. If an antivirus program scans the frequently active Workload Balancing database, it impedes or slows down the normal operation of Workload Balancing. Consequently, you must configure antivirus software running on your Workload Balancing servers to exclude specific processes and files. Citrix recommends configuring your antivirus software to exclude these folders before you initialize Workload Balancing and begin collecting data.

To configure antivirus software on the servers running Workload Balancing components:

- Exclude the following folder, which contains the Workload Balancing log:

On Windows XP and Windows Server 2003: C:\Documents and Settings\All Users\Application Data\Citrix\Workload Balancing\Data\Logfile.log

On Windows Vista and Windows Server 2008: C:\Program Data\Citrix\Workload Balancing\Data\Logfile.log.

- Exclude the SQL Server database folder. For example:

On SQL Server: C:\Program Files\Microsoft SQL Server\MSSQL\Data\

On SQL Server Express: C:\Program Files\Microsoft SQL Server\MSSQL10.SQLEXPRESS\MSSQL\Data\

These paths may vary according to your operating system and SQL Server version.

Note:

These paths and file names are for 32-bit default installations. Use the values that apply to your installation. For example, paths for 64-bit edition files might be in the C:\Program Files (x86) folder.

Configuring Workload Balancing Settings

After initialization, you can edit the Workload Balancing server connection settings and the settings Workload Balancing uses to calculate placement and optimization recommendations.

Placement and optimization settings you can perform include changing the placement strategy, configuring automatic optimizations and power management, editing performance thresholds and metric weightings, and excluding hosts. You can also set up report subscriptions that generate and email reports to you on a preset schedule automatically.

Workload Balancing settings apply collectively to all virtual machines and hosts in the pool. However, you must configure individual settings for each resource pool in your environment.

Provided the network and disk thresholds align with the hardware in your environment, you might want to consider using most of the defaults in Workload Balancing initially.

After Workload Balancing is enabled for a while, Citrix recommends evaluating your performance thresholds and determining if you need to edit them. For example, consider if you are:

- Getting optimization recommendation when they are not yet required. If this is the case, try adjusting the thresholds until Workload Balancing begins providing suitable optimization recommendations.
- Not getting recommendations when you think your network has insufficient bandwidth. If this is the case, try lowering the network critical thresholds until Workload Balancing begins providing optimization recommendations.



Before you edit your thresholds, you might find it handy to generate a [the section called “Host Health History”](#) report for each physical host in the pool.

You can use either the **Workload Balancing Configuration** properties in XenCenter or the XE commands to modify the configuration settings.

To update the credentials XenServer and the Workload Balancing server use to communicate, see [Updating Workload Balancing Credentials](#).

To display the Workload Balancing Configuration dialog box

1. In the **Resources** pane of XenCenter, select XenCenter > `resourcelessness`.
2. In the **Properties** pane, click the **WLB** tab.
3. In the **WLB** tab, click **Configure WLB**.

Adjusting the Optimization Mode

Workload Balancing makes recommendations to rebalance, or *optimize*, the virtual-machine workload in your environment based on a strategy for placement you select known as the *optimization mode*.

Workload Balancing lets you choose from two optimization modes:

- **Maximize Performance.** (Default.) Workload Balancing attempts to spread workload evenly across all physical hosts in a resource pool. The goal is to minimize CPU, memory, and network pressure for all hosts. When Maximize Performance is your placement strategy, Workload Balancing recommends optimization when a virtual machine reaches the High threshold.
- **Maximize Density.** Workload Balancing attempts to fit as many virtual machines as possible onto a physical host. The goal is to minimize the number of physical hosts that must be online.

When you select Maximize Density as your placement strategy, you can specify rules similar to the ones in Maximize Performance. However, Workload Balancing uses these rules to determine how it can pack virtual machines onto a host. When Maximize Density is your placement strategy, Workload Balancing recommends optimization when a virtual machine reaches the Critical threshold.

Workload Balancing also lets you apply these optimization modes all of the time, *Fixed*, or switch between modes for specified time periods, *Scheduled*.

Fixed

Fixed optimization modes set Workload Balancing to a specific optimization behavior - either to try to create the best performance or the highest density - at all times.

Scheduled

Scheduled optimization modes let you schedule for Workload Balancing to apply different optimization modes depending on the time of day. For example, you might want to configure Workload Balancing to optimize for performance during the day when you have users connected and then, to save energy, specify for Workload Balancing to optimize for Maximum Density at night.

When you configure Scheduled optimization modes, Workload Balancing automatically changes to the optimization mode at the beginning of the time period you specified. You can select to either configure

weekdays or weekends as a group when specifying days. For the hour, you can choose any hour and you can specify minutes in 15-minute intervals.

To set an optimization mode for all time periods

1. In the **Resources** pane of XenCenter, select XenCenter > `your-resource-pool`.
2. In the **Properties** pane, click the **WLB** tab.
3. On the **WLB** tab, click **Configure WLB**.
4. In the left pane, click **Optimization Mode**.
5. In the **Fixed** section of the **Optimization Mode** page, select one of these optimization modes:
 - **Maximize Performance**. (Default.) Attempts to spread workload evenly across all physical hosts in a resource pool. The goal is to minimize CPU, memory, and network pressure for all hosts.
 - **Maximize Density**. Attempts to fit as many virtual machines as possible onto a physical host. The goal is to minimize the number of physical hosts that must be online.

To specify times when the optimization mode will change automatically

1. In the **Resources** pane of XenCenter, select XenCenter > `your-resource-pool`.
2. In the **Properties** pane, click the **WLB** tab.
3. On the **WLB** tab, click **Configure WLB**.
4. In the left pane, click **Optimization Mode** and select **Scheduled**.
5. Click **Add New**.
6. In the **Change to** box, select one of the following modes:
 - **Maximize Performance**. Attempts to spread workload evenly across all physical hosts in a resource pool. The goal is to minimize CPU, memory, and network pressure for all hosts.
 - **Maximize Density**. Attempts to fit as many virtual machines as possible onto a physical host. The goal is to minimize the number of physical hosts that must be online.
7. Select the day of the week and the time when you want Workload Balancing to begin operating in this mode.
8. Create more scheduled mode changes (that is, "tasks") until you have the number you need. If you only schedule one task, then Workload Balancing assumes that your starting mode is Maximize Performance.
9. (Optional.) If you do not want this schedule to go into effect, clear the **Enable Task** check box.
10. Click **OK**.

To edit or delete an automatic optimization interval

1. Follow steps 1 to 4 in the previous procedure.
2. Select a scheduled task and click one of the following:
 - **Edit**. Modify the schedule or optimization mode for a task as described in step 4 to 8 of the previous procedure.
 - **Delete**. Deletes the task.

Optimizing and Managing Power Automatically

You can configure Workload Balancing to accept optimization recommendations automatically (*Automation*) and turn servers on or off automatically (*Power Management*).



Accepting Optimization Recommendations Automatically

Workload Balancing lets you configure for it to accept optimization recommendations on your behalf and perform the optimization actions it recommends automatically. This feature is required when you configure for hosts with light workloads to be powered down automatically (Maximum Density mode with Power Management).

When you enable Automation, you can specify how many minutes you want Workload Balancing to wait before it accepts (applies) its optimization recommendations.

This functionality is not enabled by default, so if want to have it enabled, you must configure it. If you do not, Workload Balancing still prompts you to accept optimization recommendations manually.

Enabling Workload Balancing Power Management

The term *power management* refers to the ability to turn the power on or off for physical hosts in a resource pool based on the pool's total workload.

Configuring power management on a host requires that:

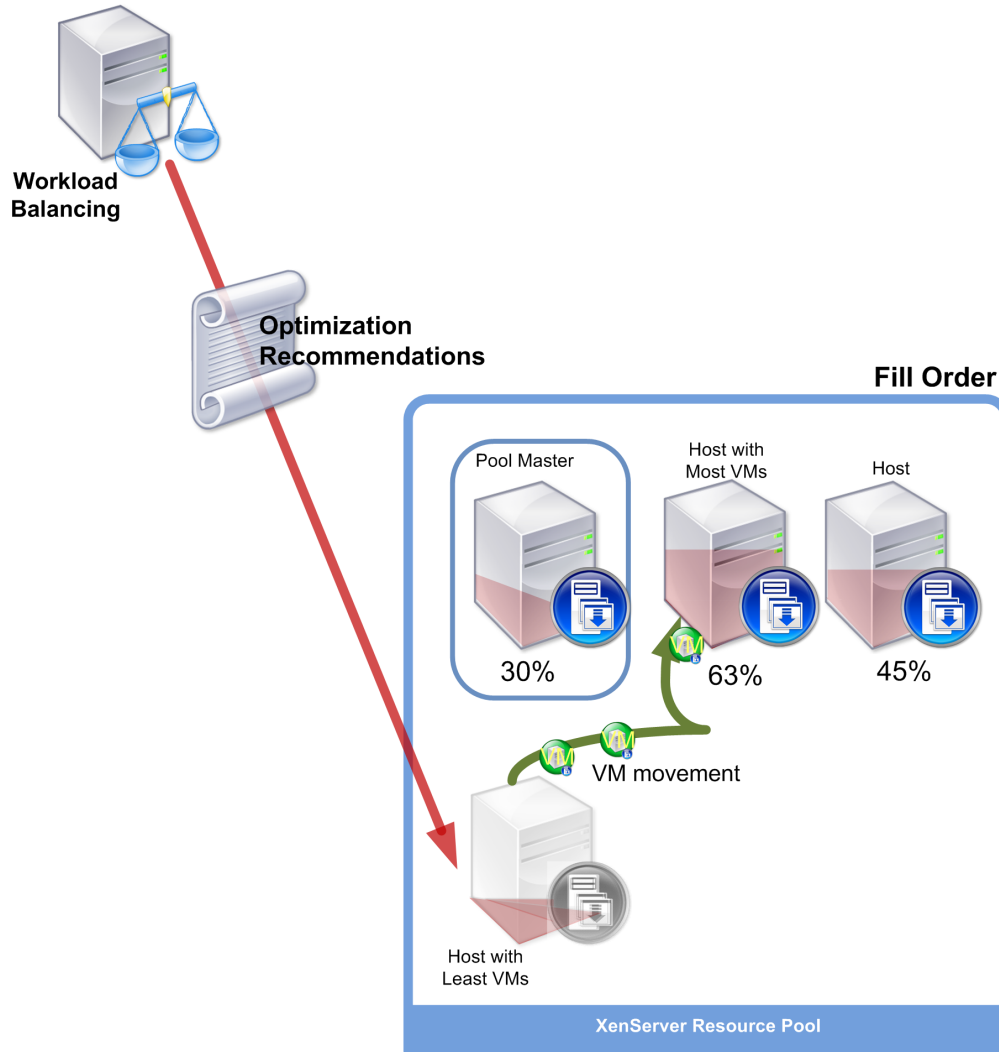
- The hardware for the host server has remote power on/off capabilities
- The Host Power On feature is configured for the host
- The pool's optimization mode is set to Maximum Density or Maximum Performance, either as a Fixed mode or a Scheduled mode
- Workload Balancing is configured to apply Optimization recommendations automatically
- Workload Balancing is configured to apply Power Management recommendations automatically
- The host has been explicitly selected as a host to be managed under Power Management

When enabled, if Workload Balancing detects unused resources, it recommends powering off hosts until it eliminates all excess capacity in the pool. If Workload Balancing detects there is not sufficient host capacity in the pool to turn off servers, it recommends leaving the servers on until the pool's workload decreases enough to power off servers. When you configure Workload Balancing to power off extra servers automatically, it applies these recommendations automatically and, consequently, behaves in the same way.

When you enable Power Management, you do so at the resource pool level. However, you can specify the individual hosts from the pool for which you want enable Power Management.

Understanding Power Management Behavior

Before Workload Balancing powers servers on or off, it selects the hosts to transfer virtual machines to (that is, to "fill") according to which hosts have the most virtual machines running. (This means that, to encourage density, it fills the "most loaded" server first.) The host with the most virtual machines is filled first (takes precedence). The pool master is filled last (lowest level of precedence) so that it is not overloaded.



If Workload Balancing detects a performance issue while the pool is in Maximum Density mode, it attempts to address the issue by recommending migrating workloads among the powered-on hosts. If Workload Balancing cannot resolve the issue using this method, it attempts to power-on a host. (Workload Balancing determines which host(s) to power-on by applying the same criteria it would if the optimization mode was set to Maximum Performance.)

If, while migrating one or more virtual machines, Workload Balancing determines that increasing capacity benefit the pool's overall performance, it turns hosts back on automatically or recommends doing so.

Important:

Workload Balancing never recommends powering on a host unless Workload Balancing powered it off.

Designing Environments for Power Management and VM Consolidation

When you are planning a XenServer implementation and you intend to configure automatic VM consolidation and power management, consider your workload design. For example, you may want to:

- **Place Different Types of Workloads in Separate Pools.** If you have an environment with distinct types of workloads (for example, user applications versus domain controllers or types of applications that perform

better with certain types of hardware), consider if you need to locate the virtual machines hosting these workloads in different pools.

Because power management and VM consolidation are managed at the pool level, you should design pools so they contain workloads that you want consolidated at the same rate, factoring in considerations such as those discussed in [the section called “Configuring Optimization Intervals, Report Subscriptions, and Data Storage”](#).

- **Exclude Hosts from Workload Balancing.** Some hosts may need to be on at all times. See [the section called “Excluding Hosts from Recommendations”](#) for more information.

To apply optimization recommendations automatically

1. In the **Resources** pane of XenCenter, select XenCenter > `your-resource-pool`.
2. In the **Properties** pane, click the **WLB** tab.
3. In the **WLB** tab, click **Configure WLB**.
4. In the left pane, click **Automation/Power Management**.
5. Select one or more of the following check boxes:
 - **Automatically apply Optimization recommendations.** When you select this option, you do not need to accept optimization recommendations manually. Workload Balancing automatically accepts optimization and placement recommendations it makes.
 - **Automatically apply Power Management recommendations.** When you select this option, it specifies for Workload Balancing to implement Power Management recommendations it makes without requiring you to accept them manually. If you want Workload Balancing to power off hosts automatically during low usage periods, select this option. To apply power recommendations automatically, Workload Balancing requires you to configure it to apply optimizations automatically.
6. Specify the number of minutes Workload Balancing waits before it applies an optimization recommendation. The default is three minutes.
7. Select the lowest level of optimization recommendation that you want Workload Balancing to apply automatically. The default is **High**.
8. Continue on to select specific servers for power management or exit the dialog.

To select servers for power management

1. In the Power Management section, select the hosts that you want Workload Balancing to power on and off automatically.

Note:

Selecting hosts for power management recommendations without selecting the Automatically apply Power Management recommendations check box results in Workload Balancing suggesting power management recommendations but not applying them automatically for you.

2. Click OK. If none of the physical servers in the resource pool support remote power management, Workload Balancing displays the message, "No hosts support Power Management."

Changing the Critical Thresholds

Workload Balancing evaluates CPU, Memory, Network Read, Network Write, Disk Read, and Disk Write utilization for physical hosts in a resource pool.



Workload Balancing determines whether to recommend relocating a workload and whether a physical host is suitable for a virtual-machine workload by evaluating:

- Whether a resource's critical threshold is met on the physical host
- (If the critical threshold is met) the importance assigned to a resource

Note:

To prevent data from appearing artificially high, Workload Balancing evaluates the daily averages for a resource and smooths utilization spikes.

When evaluating utilization, Workload Balancing compares its daily average to four thresholds: low, medium, high, and critical. After you specify (or accept the default) critical threshold, Workload Balancing sets the other thresholds relative to the critical threshold on a pool.

Default Settings for Critical Thresholds

The default settings for the Critical Thresholds are as follows:

Setting	Default
CPU Utilization	90%
Free Memory	51 MB
Network Read	26 MB per second
Network Write	26 MB per second
Disk Read	26 MB per second
Disk Write	26 MB per second

To prevent the pool master from becoming overloaded, Workload Balancing automatically sets the pool master's Critical Thresholds at lower values.

To change the critical thresholds

1. In the **Resources** pane of XenCenter, select XenCenter > `your-resource-pool`.
2. In the **Properties** pane, click the **WLB** tab.
3. In the **WLB** tab, click **Configure WLB**.
4. In the left pane, select **Critical Thresholds**.
5. In **Critical Thresholds** page, accept or enter a new value in the **Critical Thresholds** boxes.

Workload Balancing uses these thresholds when making virtual-machine placement and pool-optimization recommendations. Workload Balancing strives to keep resource utilization on a host below the critical values set.

Tuning Metric Weightings

Balancing workloads occurs if a resource's utilization is significant enough to warrant or prevent relocating a workload. For example, if you set memory as a **Less Important** factor in placement recommendations, Workload Balancing may still recommend placing virtual machines you are relocating on a server with high-memory utilization.

The effect of the weighting varies according to the placement strategy you selected. For example, if you selected **Maximum Performance** and you set **Network Writes** towards **Less Important**, if the **Network Writes** on that server exceed the critical threshold you set, Workload Balancing still makes a recommendation to place a virtual machine's workload on a server but does so with the goal of ensuring performance for the other resources.

If you selected **Maximum Density** as your placement recommendation and you specify **Network Writes** as **Less Important**, Workload Balancing will still recommend placing workloads on that host if the Network Writes exceed the critical threshold you set. However, the workloads are placed in the densest possible way.

By default, all metric weightings are set to the farthest point on the slider (**More Important**).

To edit metric weighting factors

1. In the **Resources** pane of XenCenter, select XenCenter > `your-resource-pool`.
2. In the **Properties** pane, click the **WLB** tab.
3. In the **WLB** tab, click **Configure WLB**.
4. In the left pane, select **Metric Weighting**.
5. In **Metric Weighting** page, if desired, adjust the sliders beside the individual resources.

Moving the slider towards **Less Important** indicates that ensuring virtual machines always have the highest amount of this resource available is not as vital on this resource pool.

Excluding Hosts from Recommendations

When configuring Workload Balancing, you can specify that specific physical hosts are excluded from Workload Balancing optimization and placement recommendations, including Start On placement recommendations.

Situations when you may want to exclude hosts from recommendations include when:

- You want to run the pool in Maximum Density mode and consolidate and shut down hosts, but there are specific hosts you want to exclude from this behavior.
- When two VM workloads always need to run on the same host (for example, if they have complementary applications or workloads).
- You have workloads that you do not want moved around a lot (for example, domain controllers or SQL Server).
- You want to perform maintenance on a host and you want to leave the host on the network (in the pool).
- The performance of the workload is so critical that the cost of dedicated hardware is irrelevant.
- Specific hosts are running high-priority workloads, which you do not want to prioritize using the High Availability feature.
- The hardware in the host is not optimum for the other workloads in the pool.

Regardless of whether you specify a fixed or scheduled optimization mode, hosts excluded remain excluded even when the optimization mode changes. Therefore, if you only want to prevent Workload Balancing from shutting off a host automatically, consider not enabling (or deselecting) Power Management for that host instead as described in *Optimizing and Managing Power Automatically*.

To exclude hosts from placement and optimization recommendations

1. In the **Resources** pane of XenCenter, select XenCenter > `your-resource-pool`.



2. In the **Properties** pane, click the **WLB** tab.
3. In the **WLB** tab, click **Configure WLB**.
4. In the left pane, select **Excluded Hosts**.
5. In **Excluded Hosts** page, select the hosts for which you do not want WLB to recommend alternate placements and optimizations.

Configuring Optimization Intervals, Report Subscriptions, and Data Storage

Workload Balancing supplies some advanced features for you to fine-tune your configuration, including:

Historical Data (Storage Time)

For environments using SQL Server 2005 or 2008 (not SQL Server Express), you can specify the number of weeks that Workload Balancing retains the data about your environment's workloads and optimizations. After this period, unless you move the data from the data store to another server (archive it), the data is deleted and is no longer available in reports.

For environments with SQL Express, Workload Balancing automatically grooms archived data to keep the database size below the SQL Express limit of 4 GB.

To configure the data storage period

- If you want to change the number of weeks this historical data should be stored for this resource pool, type a new value in the **Weeks** box. This option is not available if the data store is on SQL Server Express. (For SQL Express, Workload Balancing controls how large your database can become and limits it to 4GB.)

VM Optimization Criteria

The frequency of VM recommendation intervals is a product of multiple factors, including:

- How long you specify Workload Balancing waits before applying another optimization recommendation
- How many recommendations Workload Balancing must make before applying the recommendation
- The severity level a recommendation must achieve optimization
- The assertiveness with which Workload Balancing applies recommendations

Length of Time Between Optimization Recommendations After VM Moves

You can specify the number of times Workload Balancing must wait after the last time a virtual machine was moved, regardless of the cause, before it can generate another optimization recommendation that includes virtual machine.

Setting the relocation interval prevents Workload Balancing from generating optimization recommendations for artificial reasons (for example, if there was a temporary utilization spike).

The frequency for which you should set the optimization generation interval varies according to your environment. Consider these scenarios:

- If server loads and activity increase extremely quickly in your environment, you may want to decrease number of minutes for the optimization interval.

For example, if you know that activity in your environment tends to increase sharply every morning because all employees arrive to work and access their published applications exactly the same time every day (for example, 9:00 AM) and your enrolment always spikes for forty minutes, you prevent workload

balancing from generating "artificial" recommendations for this period by setting the Relocation interval to more than an hour.

- If server loads and activity increase gradually in your environment, you may want to increase the number of minutes for the optimization interval.

When Automation is configured, it is especially important to be careful when modifying the optimization interval. If an issue occurs that leads to continuous, recurring spikes, increasing the frequency (setting a lower number) can generate many recommendations and, consequently, relocations.

Note:

Setting an optimization interval does not affect how long Workload Balancing waits to factor recently rebalanced servers into the servers it recommends for Start-On placement, Resume, and Maintenance Mode.

Number of Times an Optimization Recommendation is Made

When you enable Automation, you can specify how many times you want Workload Balancing to wait before it accepts (applies) its optimization recommendations. Accepting optimization recommendations uses system resources and affects performance while Workload Balancing is relocating the virtual machines. Setting this interval prevents Workload Balancing from applying optimization recommendations too soon after the last optimization (for example, if there was a temporary utilization spike). The interval you choose to set depends on factors such as the following:

- If server loads and activity increase extremely quickly in your environment, you may want to decrease the optimization interval.
- If server loads and activity increase gradually in your environment, you may want to increase the optimization interval.

Citrix does not recommend modifying this setting unless Citrix Technical Support directs you to do so.

Setting the Minimum Optimization Severity

All optimization recommendations include a severity rating (Critical, High, Medium, Low) that indicates the importance of the recommendation. Workload Balancing bases this rating on a combination of factors including configuration options you set, such as thresholds and metric tunings; resources available for the workload; and resource-usage history. (The severity rating for a recommendation appears in the Optimization Recommendations pane on the WLB tab.)

When you configure Workload Balancing to apply optimization recommendations automatically, you can set the minimum severity level that should be associated with a recommendation before Workload Balancing automatically applies it.

Modifying the Aggressiveness Setting

The Workload Balancing Aggressiveness setting controls how strictly Workload Balancing applies optimization criteria when determining whether to apply recommendations automatically when Automation is enabled.

The conditions under which Workload Balancing applies optimizations automatically vary according to the settings you configure in the VM Recommendations Interval section of this dialog box.

Specifically, Workload Balancing looks at these settings in this dialog conjunction with the following principles:

Before applying a recommendations, the exact same recommendation has to occur x number (where x is specified by the user). Workload Balancing checks for recommendations every two minutes.

Example recommendation: [1] Move VM07 to host18 [2] Move VM07 to host20 [3] Move VM07 to host18. If 3 was specified for *x*, this is an example where Workload Balancing would *not* automatically apply the recommendation yet.

The recommendations have to be exact: notice the move to the host is not the same for all three recommendations. When you set the aggressiveness setting higher (for example, Medium), Workload Balancing will be less strict. In this case, Workload Balancing may go ahead and apply the recommendation ignoring the fact that the move to host was not the same.

To configure VM Recommendation intervals

1. In the **Resources** pane of XenCenter, select XenCenter `your-resource-pool`.
2. In the **Properties** pane, click the **WLB** tab.
3. In the **WLB** tab, click **Configure WLB**.
4. In the left pane, click **Advanced**.
5. In the **VM Recommendation Interval** section, do one or more of the following:
 - In the **Minutes** box, type a value for the number of minutes you want Workload Balancing to wait before making another optimization recommendation on a newly rebalanced server.
 - In the **Recommendations** box, type a value for the number of optimization recommendations you want Workload Balancing to make before it applies an optimization recommendation automatically.
 - Select a minimum severity level before optimizations are applied automatically.
 - Modify how aggressively Workload Balancing applies optimization recommendations when it is running in Automated mode. Increasing the aggressiveness level reduces constraints on the consistency of recommendations before automatically applying them. The Aggressiveness setting directly complements the **Number of Optimization Recommendations** setting.

Note:

If you specify "1" for the value in the **Number of Optimization Recommendations** setting, the Aggressiveness setting is not relevant.

Receiving Reports by Email Automatically (Report Subscriptions)

Workload Balancing lets you subscribe to reports so that they are delivered automatically by email on a schedule that you configure. This feature is known as *Report Subscriptions*.

Configuring report subscriptions requires two tasks:

1. Enabling Report Subscriptions by configuring the name of the email server.
2. Specifying the report to which you want to subscribe, as described in To subscribe to a Workload Balancing report.

Instead of using the default report delivery mechanism included with the report subscriptions feature, you can configure for Workload Balancing to deliver reports using SQL Server Reporting Services. In large environments, especially ones with existing SQL Server Reporting Services, it may be advantageous to configure Reporting Services as the report delivery mechanism. Environments with large amounts of Workload Balancing data achieve better reporting performance and faster report generation times with Reporting Services.

Note:



If you want to use the default report delivery mechanism, no additional component installations or configurations are required.

To configure report subscriptions

1. In the **Resources** pane of XenCenter, select XenCenter > `your-resource-pool`.
2. In the **Properties** pane, click the **WLB** tab.
3. In the **WLB** tab, click **Configure WLB**.
4. In the left pane, click **Advanced**.
5. In the **Report Subscriptions** section, in the SMTP box, type the IP address or fully qualified domain name (FQDN) of the email server that Workload Balancing will use to email the reports.
6. (Optional.) If you want to use Reporting Services as your report delivery mechanism, select the **Use Reporting Services for Report Delivery** check box.

Choosing an Optimal Server for VM Initial Placement, Migrate, and Resume

When Workload Balancing is enabled and you restart a virtual machine that is offline, XenCenter provides recommendations to help you determine the *optimal* physical host in the resource pool on which to start the virtual machine. Workload Balancing makes these placement recommendations by using performance metrics it previously gathered for that virtual machine and the physical hosts in the resource pool. Likewise, when Workload Balancing is enabled, if you migrate a virtual machine to another host, XenCenter recommends servers to which you can move that host. This Workload Balancing enhancement is also available for the Initial (Start On) Placement and Resume features.

In general, Workload Balancing functions more effectively and makes better, less frequent optimization recommendations if you start (Initial Start On Placement and Resume) virtual machines on the servers it recommends.

When you use these features with Workload Balancing enabled, host recommendations appear as star ratings beside the name of the physical host. Five empty stars indicates the lowest-rated (least optimal) server. When it is not possible to start or move a virtual machine to a host, an (X) appears beside the host name with the reason.

To start a virtual machine on the optimal server

1. In the **Resources** pane of XenCenter, select the virtual machine you want to start.
2. From the **VM** menu, select **Start on Server** and then select one of the following:
 - **Optimal Server.** The optimal server is the physical host that is best suited to the resource demands of the virtual machine you are starting. Workload Balancing determines the optimal server based on its historical records of performance metrics and your placement strategy. The optimal server is the server with the most stars.
 - One of the servers with star ratings listed under the Optimal Server command. Five stars indicates the most-recommended (optimal) server and five empty stars indicates the least-recommended server.

To resume a virtual machine on the optimal server

1. In the Resources pane of XenCenter, select the suspended virtual machine you want to resume.
2. From the VM menu, select Resume on Server and then select one of the following:
 - **Optimal Server.** The optimal server is the physical host that is best suited to the resource demands of the virtual machine you are starting. Workload Balancing determines the optimal server based on



its historical records of performance metrics and your placement strategy. The optimal server is the server with the most stars.

- One of the servers with star ratings listed under the Optimal Server command. Five stars indicates the most-recommended (optimal) server and five empty stars indicates the least-recommended server.

Accepting Optimization Recommendations

Workload Balancing provides recommendations about ways you can move virtual machines to optimize your environment. Optimization recommendations appear in the **WLB** tab in XenCenter. Optimization recommendations are based on the:

- Placement strategy you select (that is, the placement optimization mode), as described in [the section called “Adjusting the Optimization Mode”](#)
- Performance metrics for resources such as a physical host's CPU, memory, network, and disk utilization
- The role of the host in the resource pool. When making placement recommendations, Workload Balancing considers only the pool master if no other virtual machine can accept the workload. (Likewise, when a pool is operating in Maximum Density mode, Workload Balancing considers the pool master last when determining the order in which to fill hosts with virtual machines.)

The optimization recommendations display the name of the virtual machine that Workload Balancing recommends relocating, the host it currently resides on, and the host Workload Balancing recommends as the machine's new location. The optimization recommendations also display the reason Workload Balancing recommends moving the virtual machine (for example, "CPU" to improve CPU utilization).

After you accept an optimization recommendation, XenServer relocates all virtual machines listed as recommended for optimization.

Tip:

You can find out the optimization mode for a resource pool by selecting the pool in XenCenter and checking the **Configuration** section of the **WLB** tab.

To accept an optimization recommendation

1. In the **Resources** pane of XenCenter, select the resource pool for which you want to display recommendations.
2. In the **Properties** pane, click the **WLB** tab. If there are any recommended optimizations for any virtual machines on the selected resource pool, they display on the **WLB** tab.
3. To accept the recommendations, click **Apply Recommendations**. XenServer begins moving all virtual machines listed in the **Optimization Recommendations** section to their recommended servers.

After you click **Apply Recommendations**, XenCenter automatically displays the **Logs** tab so you can see the progress of the virtual machine migration.

Administering Workload Balancing

Some administrative tasks you may want to perform on Workload Balancing include:

- [Disabling Workload Balancing](#) on a pool, either permanently or temporarily
- [the section called “Reconfiguring a Pool to Use Another WLB Server”](#) server that a pool uses
- [Changing the credentials Workload Balancing](#) or XenServer use to communicate

- [Uninstalling Workload Balancing](#)

Disabling Workload Balancing

You can disable Workload Balancing for a resource pool, either temporarily or permanently:

- **Temporarily.** Disabling Workload Balancing temporarily stops XenCenter from displaying recommendations for the specified resource pool. When you disable Workload Balancing temporarily, data collection stops for that resource pool.
- **Permanently.** Disabling Workload Balancing permanently deletes information about the specified resource pool from the data store and stops data collection for that pool.

To disable Workload Balancing on a pool

1. In the **Resource** pane of XenCenter, select the resource pool for which you want to disable Workload Balancing.
2. In the WLB tab, click **Disable WLB**. The **Disable Workload Balancing** dialog box appears.
3. Click **Yes** to disable Workload Balancing for the pool. Important: If you want to disable Workload Balancing permanently for this resource pool, click the **Remove all resource pool information from the Workload Balancing Server** check box.

XenServer disables Workload Balancing for the resource pool, either temporarily or permanently depending on your selections.

- If you disabled Workload Balancing temporarily on a resource pool, to reenable Workload Balancing, click **Enable WLB** in the **WLB** tab.
- If you disabled Workload Balancing permanently on a resource pool, to reenable it, you must reinitialize it. For information, see [Initializing Workload Balancing](#).

Reconfiguring a Pool to Use Another WLB Server

You can reconfigure a resource pool to use a different Workload Balancing server.

However, to prevent old data collectors from remaining inadvertently configured and running against a pool, you must disable Workload Balancing permanently for that resource pool before pointing the pool to the other data collector.

After disabling Workload Balancing, you can re-initialize the pool and specify the name of the new Workload Balancing server.

To use a different Workload Balancing server

1. On the resource pool you want to point to a different Workload Balancing server, disable **Workload Balancing**. Specify to delete the pool information when you disable it. For instructions, see [the section called "Disabling Workload Balancing"](#).
2. In the **WLB** tab, click **Enable WLB**. The **Connect to WLB Server** dialog appears.
3. In the **Address** box, type the IP address or NetBIOS name of the new Workload Balancing server. You can also enter a fully qualified domain name (FQDN).
4. Reinitialize the resource pool and specify the new server's IP address in the **Connect to WLB Server** dialog. If either the Workload Balancing credentials have changed as a result of specifying a different WLB server, you must also update these credentials.

Note:



You must enter all the information that you would normally enter when you initially configure a resource pool for use with Workload Balancing. For information, see [the section called “Initializing Workload Balancing”](#) or [the section called “Updating Workload Balancing Credentials”](#).

Updating Workload Balancing Credentials

After initial configuration, if you need to update the credentials XenServer and the Workload Balancing server use to communicate, you can do so without reconfiguring all settings by using the Update Workload Balancing Credentials feature.

Situations when you may want to use this feature include:

- If you need to change the user account XenServer uses to communicate with Workload Balancing
- If you receive an error message that the Workload Balancing credentials are no longer valid
- If the service is unavailable

If you want to modify settings for thresholds and change the priority given to specific resources, see [the section called “Configuring Workload Balancing Settings”](#).

1. Temporarily disable Workload Balancing:

- In the **Resources** pane of XenCenter, select `XenCenter > your-resource-pool`.
- In the **Properties** pane, click the **WLB** tab.
- In the **WLB** tab, click **Disable WLB**.

The **Disable Workload Balancing** dialog box appears.

- Click **Yes**.

Note:

Do not select the **Remove all resource pool information from the Workload Balancing server** check box.

2. Reenable Workload Balancing and specify the new credentials:

- After the progress bar completes, click **Enable WLB**.

The **Connect to WLB Server** dialog box appears.

- Click **Update Credentials**.
- In the **Server Address** section, modify the following as desired:
 - a. In the **Address** box, type the IP address or NetBIOS name of the Workload Balancing server. You can also enter a fully qualified domain name (FQDN).
 - b. (Optional.) Edit the port number if you want XenServer to connect to Workload Balancing using a different port. Entering a new port number here sets a different communications port on the Workload Balancing server.

By default, XenServer connects to Workload Balancing (specifically the Web Service Host service) on port 8012.

Note:

Do not edit this port number unless you have changed it during Workload Balancing Setup. The port number value specified during Setup and in the **Workload Balancing Configuration** dialog must match.



- c. In the **WLB Server Credentials** section, enter the user name (for example, workloadbalancing_user) and password the computers running XenServer will use to connect to the Workload Balancing server.

This must be the account or group that was configured during the installation of the Workload Balancing Server. For information, see [the section called "Authorization for Workload Balancing"](#).

- d. In the **XenServer Credentials** section, enter the user name and password for the pool you are configuring (typically the password for the pool master). Workload Balancing will use these credentials to connect to the computers running XenServer in that pool.

To use the credentials with which you are currently logged into XenServer, select the **Use the current XenCenter credentials** check box.

Uninstalling Workload Balancing

Citrix recommends uninstalling Workload Balancing from the **Control Panel** in Windows.

When you uninstall Workload Balancing, only the Workload Balancing software is removed from the Workload Balancing server. The data store remains on the system running SQL Server. To remove a Workload Balancing data store, you must use the SQL Server Management Studio (SQL Server 2005 and SQL Server 2008).

If you want to uninstall both Workload Balancing and SQL Server from your computer, uninstall Workload Balancing first and then delete the database using the SQL Server Management Studio.

The data directory, usually located at C:\Documents and Settings\All Users\Application Data\Citrix\Workload Balancing\Data, is not removed when you uninstall Workload Balancing. You can remove the contents of the data directory manually.

Customizing Workload Balancing

Workload Balancing provides several methods of customization:

- **Database Views Documentation.** You can use these views to create custom queries on your data without having to fully understand the Workload Balancing database schema. You can use them as a guideline for creating tables if you want to import the files into your own database or archiving and data warehousing.

The database-views documentation provides three tables with information about Host Metric History, Host Virtual Machine History, and Virtual Machine Metric History. See Citrix Knowledge Center article, [Citrix XenServer Workload Balancing SQL Views], [CTX121348](#).

- **Command lines for scripting.** See the Workload Balancing section of the [XenServer Administrator's Guide].
- **Host Power On scripting support.** You can also customize Workload Balancing (indirectly) through the Host Power On scripting. How to create a custom script for Host Power On is documented in the [XenServer Administrator's Guide].

Entering Maintenance Mode with Workload Balancing Enabled

When Workload Balancing is enabled, if you take a physical host offline for maintenance (that is, suspend a server by entering Maintenance Mode), XenServer automatically migrates the virtual machines running on that host to their optimal servers when available. XenServer migrates them based on Workload Balancing recommendations (performance data, your placement strategy, and performance thresholds).



If an optimal server is not available, the words **Click here to suspend the VM** appear in the **Enter Maintenance Mode** dialog box. In this case, Workload Balancing does not recommend a placement because no host has sufficient resources to run this virtual machine. You can either suspend this virtual machine or exit Maintenance Mode and suspend a virtual machine on another host in the same pool. Then, if you reenter the **Enter Maintenance Mode** dialog box, Workload Balancing might be able to list a host that is a suitable candidate for migration.

Note:

When you take a server offline for maintenance and Workload Balancing is enabled, the words "Workload Balancing" appear in the upper-right corner of the **Enter Maintenance Mode** dialog box.

To enter maintenance mode with Workload Balancing enabled

1. In the **Resources** pane of XenCenter, select the physical host that you want to take offline. From the **Server** menu, select **Enter Maintenance Mode**.
2. In the **Enter Maintenance Mode** dialog box, click **Enter maintenance mode**. The virtual machines running on the server are automatically migrated to the optimal host based on Workload Balancing's performance data, your placement strategy, and performance thresholds.

To take the server out of maintenance mode, right-click the server and select **Exit Maintenance Mode**. When you remove a server from maintenance mode, XenServer automatically restores that server's original virtual machines to that server.

Working with Workload Balancing Reports

This topic provides general information about Workload Balancing historical reports and an overview of where to find additional information about these reports.

To generate a Workload Balancing report, you must have installed the Workload Balancing component, registered at least one resource pool with Workload Balancing, and configured Workload Balancing on at least one resource pool.

Introduction

Workload Balancing provides reporting on three types of objects: physical hosts, resource pools, and virtual machines. At a high level, Workload Balancing provides two types of reports:

- Historical reports that display information by date
- "Roll up" style reports

Workload Balancing provides some reports for auditing purposes, so you can determine, for example, the number of times a virtual machine moved.

Subscribing to Workload Balancing Reports

By subscribing to Workload Balancing reports, you can schedule for Workload Balancing to deliver individual reports with the options you want enabled on a preset schedule and have the reports delivered to your email account.

For example, to subscribe to a report, you generate a report that specifies the objects (hosts, pools, and so on) and settings you want in the subscribed report and then subscribe to it. Then, you specify the time



at which you want to generate the report, the days on which you want to generate it, and when you want to start and stop automatic report delivery.

Using Workload Balancing Reports for Tasks

The Workload Balancing reports can help you perform capacity planning, determine virtual server health, and evaluate the effectiveness of your configured threshold levels.

Evaluating the Effectiveness of Your Optimization Thresholds

You can use the Pool Health report to evaluate the effectiveness of your optimization thresholds. While Workload Balancing provides default threshold settings, you might need to adjust these defaults for them to provide value in your environment. If you don't have the optimization thresholds adjusted to the correct level for your environment, Workload Balancing recommendations might not be appropriate for your environment.

Generating and Managing Workload Balancing Reports

This topic explains how to generate, navigate, print, and export Workload Balancing reports.

To generate a Workload Balancing report

1. In XenCenter, from the **Pool** menu, select **View Workload Reports**.
2. From the Workload Reports screen, select a report from the left pane.
3. Select the **Start Date** and the **End Date** for the reporting period. Depending on the report you select, you might need to specify a host in the Host list box.
4. Click **Run Report**. The report displays in the report window.

To subscribe to a Workload Balancing report

Before you can subscribe to a report, you must enable the report subscription functionality as described in [the section called “Configuring Optimization Intervals, Report Subscriptions, and Data Storage”](#).

1. After generating a report, as described in the previous procedure, click **Subscribe**.
2. In the **New Subscription for a Report** dialog box, name the subscription.
3. Configure the following fields, as you require:
 - **To**. Use semi-colons to separate multiple email addresses.
 - **From**. Enter email address you want to appear as the sender.
 - **BCC**. Sends a copy of the report to somebody without their name appearing in the email's receiver list of report subscribers.
 - **Reply-To**. Enter the name of the person to whom report subscribers should direct their inquiries about the report. For example, if somebody has a question about a data point, you might want them to have them email the administrator for that resource pool when they click **Reply**. This is a required field.
 - **Subject**. Enter the text you want to appear in the Subject line of the email.
 - **Format**. You can choose from one of three report delivery formats: PDF, JPEG, and Excel.
 - **Comments**. Workload Balancing add the text you enter to the email that delivers the report.
4. In the **Schedule Options** section, create the schedule for delivering the reports and click **OK**:
 - Configure the duration of the time interval on which the report's data is based. For example, if you want it to always generate the most recent week's data, then select **Last Week**.
 - **Run at**. The time at which you want the report to generate.

- **Deliver On.** The day of the week you want Workload Balancing to generate and deliver the report. The schedule you are creating does not become active until you set this field to a day. Leaving this field at **None** means that no reports will be generated for this subscription.
- **Starting.** The date you want Workload Balancing to begin generating the report.
- **Ending.** The date you want Workload Balancing to stop generating the report.

After you create a subscription, a folder icon appears under the report name in the left pane.


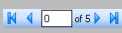





To cancel a report subscription

1. In XenCenter, from the **Pool** menu, select **View Workload Reports**.
2. From the Workload Reports screen, select the report for which you enabled the subscription from the left pane.
3. Expand the node for the report, then expand the **Subscriptions** folder.
4. Select the subscription you want to delete.
5. Confirm it is the correct subscription by reviewing the subscription contents in the right pane.
6. Click **Delete**.

To navigate in a Workload Balancing Report

After generating a report, you can use the toolbar buttons in the report to navigate and perform certain tasks. To display the name of a toolbar button, hold your mouse over toolbar icon.

Table 3. Report Toolbar Buttons

	Document Map. Lets you display a document map that helps you navigate through long reports.
	Page Forward/Back. Lets you move one page ahead or back in the report.
	Back to Parent Report. Lets you return to the parent report when working with drill-through reports.
	Stop Rendering. Cancels the report generation.
	Refresh. Lets you refresh the report display.
	Print. Lets you print a report and specify general printing options, such as the printer, the number of pages, and the number of copies.
	Print Layout. Lets you display a preview of the report before you print it.
	Page Setup. Lets you specify printing options such as the paper size, page orientation, and margins.

	Export. Lets you export the report as an Acrobat (.PDF) file or as an Excel file with a .XLS extension.
	Find. Lets you search for a word in a report, such as the name of a virtual machine.

To print a Workload Balancing report

Citrix recommends printing Workload Balancing reports in Landscape orientation.

1. After generating the report, click



Page Setup.

Page Setup also lets you control the margins and paper size.

2. In the **Page Setup** dialog, select **Landscape** and click **OK**.
3. (Optional.) If you want to preview the print job, click



Print Layout.

4. Click



Print.

To export a Workload Balancing report

You can export a report in Microsoft Excel and Adobe Acrobat (PDF) formats.

After generating the report, click



Export and select one of the following:

- **Excel**
- **Acrobat (PDF) file**

Displaying Workload Balancing Reports

The **Workload Reports** window lets you generate reports for physical hosts, resource pools, and virtual machines.







Report Generation Features

To generate a report, select a report type, the date range, the host (if applicable), and click **Run Report**. For more detail, see [the section called “To generate a Workload Balancing report”](#).

Toolbar Buttons

The following toolbar buttons in the Workload Reports window become available after you generate a report. To display the name of a toolbar button, hold your mouse over toolbar icon.

Table 4. Report Toolbar Buttons

	Document Map. Lets you display a document map that helps you navigate through long reports.
	Page Forward/Back. Lets you move one page ahead or back in the report.
	Back to Parent Report. Lets you return to the parent report when working with drill-through reports.
	Stop Rendering. Cancels the report generation.
	Refresh. Lets you refresh the report display.
	Print. Lets you print a report and specify general printing options, such as the printer, the number of pages, and the number of copies.
	Print Layout. Lets you display a preview of the report before you print it.
	Page Setup. Lets you specify printing options such as the paper size, page orientation, and margins.
	Export. Lets you export the report as an Acrobat (.PDF) file or as an Excel file with a .XLS extension.
	Find. Lets you search for a word in a report, such as the name of a virtual machine.

Workload Balancing Report Glossary

This topic provides information about the following Workload Balancing reports.

Host Health History

This report displays the performance of resources (CPU, memory, network reads, and network writes) on specific host in relation to threshold values.

The colored lines (red, green, yellow) represent your threshold values. You can use this report with the Pool Health report for a host to determine how a particular host's performance might be affecting overall pool health. When you are editing the performance thresholds, you can use this report for insight into host performance.

You can display resource utilization as a daily or hourly average. The hourly average lets you see the busiest hours of the day, averaged, for the time period.

To view report data grouped by hour, expand **+ Click to view report data grouped by house for the time period** under the **Host Health History** title bar.



Workload Balancing displays the average for each hour for the time period you set. The data point is based on a utilization average for that hour for all days in the time period. For example, in a report for May 1, 2009 to May 15, 2009, the Average CPU Usage data point represents the resource utilization of all fifteen days at 12:00 hours combined together as an average. That is, if CPU utilization was 82% at 12PM on May 1st, 88% at 12PM on May 2nd, and 75% on all other days, the average displayed for 12PM is 76.3%.

Note:

Workload Balancing smooths spikes and peaks so data does not appear artificially high.

Pool Optimization Performance History

The optimization performance report displays optimization events (that is, when you optimized a resource pool) against that pool's average resource usage. Specifically, it displays resource usage for CPU, memory, network reads, and network writes.

The dotted line represents the average usage across the pool over the period of days you select. A blue bar indicates the day on which you optimized the pool.

This report can help you determine if Workload Balancing is working successfully in your environment. You can use this report to see what led up to optimization events (that is, the resource usage before Workload Balancing recommended optimizing).

This report displays average resource usage for the day; it does not display the peak utilization, such as when the system is stressed. You can also use this report to see how a resource pool is performing if Workload Balancing is not making optimization recommendations.

In general, resource usage should decline or be steady after an optimization event. If you do not see improved resource usage after optimization, consider readjusting threshold values. Also, consider whether or not the resource pool has too many virtual machines and whether or not new virtual machines were added or removed during the time frame you specified.

Pool Audit Log History

This report displays the contents of the XenServer Audit Log, a XenServer feature designed to log attempts to perform unauthorized actions and select authorized actions, including import/export, host and pool back ups, and guest and host console access. For information about the Audit Log feature, see the audit log documentation in the XenCenter online help or the [XenServer Administrator's Guide].

To run the audit log report, the Audit Logging feature must be enabled, which it is by default.

This report displays the following:

- **Time.** The time XenServer recorded the user's action.
- **Host Name.** The name of the physical server on which the user performed the action (where the action occurred).
- **User Name.** The name of the person who created the session in which the action was performed. In some cases, this may be the User ID.
- **Access.** Whether or not the user had permission to perform the action.
- **Call Type.** The Call Type indicates whether the user connected to the session over the Internet or locally. For example, if the user connected to XenCenter over the Internet.
- **Event Object.** The object that was the subject of the action (for example, a VM).
- **Object Name.** The name of the object (for example, the name of the VM).

- **Event Action.** The action that occurred. For definitions of these actions, see [the section called “Audit Log Event Names”](#).
- **Succeeded.** This provides that status of the action (that is, whether or not it was successful).

Audit Log Event Names

The following table defines some events that appear frequently in the XenServer Audit Log report. For additional event definitions, see the events section of the *Citrix XenServer Management API*

Event Name	Definition
task.create	Created a new task object.
task.destroy	Destroyed the task object.
SR.destroy	Destroyed the storage repository.
VM.start	Started a virtual machine.
http/get_audit_log	Gets Audit Log.
pool.retrieve_wlb_recommendations	Somebody retrieved the Workload Balancing recommendations for this pool.

Pool Health

The pool health report displays the percentage of time a resource pool and its hosts spent in four different threshold ranges: Critical, High, Medium, and Low. You can use the Pool Health report to evaluate the effectiveness of your performance thresholds.

A few points about interpreting this report:

- Resource utilization in the Average Medium Threshold (blue) is the optimum resource utilization regardless of the placement strategy you selected. Likewise, the blue section on the pie chart indicates the amount of time that host used resources optimally.
- Resource utilization in the Average Low Threshold Percent (green) is not necessarily positive. Whether Low resource utilization is positive depends on your placement strategy. For example, if your placement strategy is Maximum Density and most of the time your resource usage was green, Workload Balancing might not be fitting the maximum number of virtual machines possible on that host or pool. If this is the case, you should adjust your performance threshold values until the majority of your resource utilization falls into the Average Medium (blue) threshold range.
- Resource utilization in the Average Critical Threshold Percent (red) indicates the amount of time average resource utilization met or exceeded the Critical threshold value.

If you double-click on a pie chart for a host's resource usage, XenCenter displays the Host Health History report for that resource (for example, CPU) on that host. Clicking the **Back to Parent Report** toolbar button returns you to the Pool Health history report.

If you find the majority of your report results are not in the Average Medium Threshold range, you probably need to adjust the Critical threshold for this pool. While Workload Balancing provides default threshold settings, these defaults are not effective in all environments. If you do not have the thresholds adjusted to the correct level for your environment, Workload Balancing's optimization and placement recommendations might not be appropriate. For more information, see [the section called “Changing the Critical Thresholds”](#).

Note:

The High, Medium, and Low threshold ranges are based on the Critical threshold value you set when you initialized Workload Balancing.

Pool Health History

This report provides a line graph of resource utilization on all physical hosts in a pool over time. It lets you see the trend of resource utilization - if it tends to be increasing in relation to your thresholds (Critical, High, Medium, and Low). You can evaluate the effectiveness of your performance thresholds by monitoring trends of the data points in this report.

Workload Balancing extrapolates the threshold ranges from the values you set for the Critical thresholds when you initialized Workload Balancing. Although similar to the Pool Health report, the Pool Health History report displays the average utilization for a resource on a specific date rather than the amount of time overall the resource spent in a threshold.

With the exception of the Average Free Memory graph, the data points should never average above the Critical threshold line (red). For the Average Free Memory graph, the data points should never average *below* the Critical threshold line (which is at the bottom of the graph). Because this graph displays *free* memory, the Critical threshold is a low value, unlike the other resources.

A few points about interpreting this report:

- When the Average Usage line in the chart approaches the Average Medium Threshold (blue) line, it indicates the pool's resource utilization is optimum regardless of the placement strategy configured.
- Resource utilization approaching the Average Low Threshold (green) is not necessarily positive. Whether Low resource utilization is positive depends on your placement strategy. For example, if your placement strategy is Maximum Density and most days the Average Usage line is at or below the green line, Workload Balancing might not be placing virtual machines as densely as possible on that pool. If this is the case, you should adjust the pool's Critical threshold values until the majority of its resource utilization falls into the Average Medium (blue) threshold range.
- When the Average Usage line intersects with the Average Critical Threshold Percent (red), this indicates the days when the average resource utilization met or exceeded the Critical threshold value for that resource.

If you find the data points in the majority of your graphs are not in the Average Medium Threshold range, but you are satisfied with the performance of this pool, you might need to adjust the Critical threshold for this pool. For more information, see [the section called "Changing the Critical Thresholds"](#).

Pool Optimization History

The Pool Optimization History report provides chronological visibility into Workload Balancing optimization activity.

Optimization activity is summarized graphically and in a table. Drilling into a date field within the table displays detailed information for each pool optimization performed for that day.

This report lets you see the following information:

- **VM Name.** The name of the virtual machine that Workload Balancing optimized.
- **Reason.** The reason for the optimization.
- **Status.** If the optimization was successful.
- **From Host.** The physical server where the virtual machine was originally hosted.



- **To Host.** The physical server where the virtual machine was moved.
- **Time.** The time when the optimization occurred.

Virtual Machine Motion History

This line graph displays the number of times virtual machines moved on a resource pool over a period of time. It indicates if a move resulted from an optimization recommendation and to which host the virtual machine moved. This report also indicates the reason for the optimization. You can use this report to audit the number of moves on a pool.

Some points about interpreting this report:

- The numbers on the left side of the chart correspond with the number of moves possible, which is based on how many virtual machines are in a resource pool.
- You can look at details of the moves on a specific date by expanding the + sign in the **Date** section of the report.

Virtual Machine Performance History

This report displays performance data for each virtual machine on a specific host for a time period you specify. Workload Balancing bases the performance data on the amount of virtual resources allocated for the virtual machine. For example, if the Average CPU Usage for your virtual machine is 67%, this means that your virtual machine was using, on average, 67% of its virtual CPU for the period you specified.

The initial view of the report displays an average value for resource utilization over the period you specified.

Expanding the + sign displays line graphs for individual resources. You can use these graphs to see trends in resource utilization over time.

This report displays data for CPU Usage, Free Memory, Network Reads/Writes, and Disk Reads/Writes.

Backup and recovery

This chapter presents the functionality designed to give you the best chance to recover your XenServer from a catastrophic failure of hardware or software, from lightweight metadata backups to full VM backups and portable SRs.

Backups

Citrix recommends that you frequently perform as many of the following backup procedures as possible to recover from possible server and/or software failure.

To backup pool metadata

1. Run the command:

```
xe pool-dump-database file-name=<backup>
```

2. Run the command:

```
xe pool-restore-database file-name=<backup> dry-run=true
```

This command checks that the target machine has an appropriate number of appropriately named NICs, which is required for the backup to succeed.

To backup host configuration and software

- Run the command:

```
xe host-backup host=<host> file-name=<hostbackup>
```

Note:

- Do not create the backup in the control domain.
- This procedure may create a large backup file.
- To complete a restore you have to reboot to the original install CD.
- This data can only be restored to the original machine.

To backup a VM

1. Ensure that the VM to be backed up is offline.
2. Run the command:

```
xe vm-export vm=<vm_uuid> filename=<backup>
```

Note:

This backup also backs up all of the VM's data. When importing a VM, you can specify the storage mechanism to use for the backed up data.

Warning:

Because this process backs up all of the VM data, it can take some time to complete.

To backup VM metadata only

- Run the command:

```
xe vm-export vm=<vm_uuid> filename=<backup> --metadata
```

Full metadata backup and disaster recovery (DR)

This section introduces the concept of Portable Storage Repositories (Portable SRs), and explains how they work and how to use them as part of a DR strategy.

DR and metadata backup overview

XenServer 5.6 introduces the concept of Portable SRs. Portable SRs contain all of the information necessary to recreate all the Virtual Machines (VMs) with Virtual Disk Images (VDIs) stored on the SR after re-attaching the SR to a different host or pool. Portable SRs can be used when regular maintenance or disaster recovery requires manually moving a SR between pools or standalone hosts.

Using portable SRs has similar constraints to XenMotion as both cases result in VMs being moved between hosts. To use portable SRs:

- The source and destination hosts must have the same CPU type and networking configuration. The destination host must have a network of the same name as the one of the source host.
- The SR media itself, such as a LUN for iSCSI and FibreChannel SRs, must be able to be moved, re-mapped, or replicated between the source and destination hosts
- If using tiered storage, where a VM has VDIs on multiple SRs, all required SRs must be moved to the destination host or pool
- Any configuration data required to connect the SR on the destination host or pool, such as the target IP address, target IQN, and LUN SCSI ID for iSCSI SRs, and the LUN SCSI ID for FibreChannel SRs, must be maintained manually
- The backup metadata option must be configured for the desired SR

Note:

When moving portable SRs between pools the source and destination pools are not required to have the same number of hosts. Moving portable SRs between pools and standalone hosts is also supported provided the above constraints are met.

Portable SRs work by creating a dedicated metadata VDI within the specified SR. The metadata VDI is used to store copies of the pool or host database as well as the metadata describing the configuration of each VM. As a result the SR becomes fully self-contained, or portable, allowing it to be detached from one host and attached to another as a new SR. Once the SR is attached a restore process is used to recreate all of the VMs on the SR from the metadata VDI. For disaster recovery the metadata backup can be scheduled to run regularly to ensure the metadata SR is current.

The metadata backup and restore feature works at the command-line level and the same functionality is also supported in xsconsole. It is not currently available through XenCenter.

Backup and restore using xsconsole

When a metadata backup is first taken, a special backup VDI is created on a SR. This VDI has an `ext3` filesystem that stores the following versioned backups:

- A full pool-database backup.
- Individual VM metadata backups, partitioned by the SRs in which the VM has disks.
- SR-level metadata which can be used to recreate the SR description when the storage is reattached.

On the XenServer host menu-driven text console, under the **Backup, Update and Restore** menu there are options which provide a more user-friendly interface to these scripts. The operations should only be performed on the pool master. You can use these menu options to perform 3 operations:

- Schedule a regular metadata backup to the default pool SR, either daily, weekly or monthly. This will regularly rotate metadata backups and ensure that the latest metadata is present for that SR without any user intervention being required.
- Trigger an immediate metadata backup to the SR of your choice. This will create a backup VDI if necessary, and attach it to the host and backup all the metadata to that SR. Use this option if you have made some changes which you want to see reflected in the backup immediately.
- Perform a metadata restoration operation. This will prompt you to choose an SR to restore from, and then the option of restoring only VM records associated with that SR, or all the VM records found (potentially from other SRs which were present at the time of the backup). There is also a **dry run** option to see which VMs would be imported, but not actually perform the operation.

For automating this scripting, there are some commands in the control domain which provide an interface to metadata backup and restore at a lower level than the menu options:

- **xe-backup-metadata** provides an interface to create the backup VDIs (with the `-c` flag), and also to attach the metadata backup and examine its contents.
- **xe-restore-metadata** can be used to probe for a backup VDI on a newly attached SR, and also selectively reimport VM metadata to recreate the associations between VMs and their disks.

Full usage information for both scripts can be obtained by running them in the control domain using the `-h` flag. One particularly useful invocation mode is **xe-backup-metadata -d** which mounts the backup VDI into dom0, and drops into a sub-shell with the backup directory so it can be examined.

Moving SRs between hosts and Pools

The metadata backup and restore options can be run as scripts in the control domain or through the **Backup, Restore, and Update** menu option in the xsconsole. All other actions, such as detaching the SR from the source host and attaching it to the destination host, can be performed using XenCenter, the menu-based xsconsole, or the xe CLI. This example uses a combination of XenCenter and xsconsole.

To create and move a portable SR using the xsconsole and XenCenter

1. On the source host or pool, in xsconsole, select the **Backup, Restore, and Update** menu option, select the **Backup Virtual Machine Metadata** option, and then select the desired SR.
2. In XenCenter, select the source host or pool and shutdown all running VMs with VDIs on the SR to be moved.
3. In the tree view select the SR to be moved and select **Storage > Detach Storage Repository**. The **Detach Storage Repository** menu option will not be displayed if there are running VMs with VDIs on the selected SR. After being detached the SR will be displayed in a grayed-out state.

Warning:

Do not complete this step unless you have created a backup VDI in step 1.

4. Select **Storage > Forget Storage Repository** to remove the SR record from the host or pool.
5. Select the destination host in the tree view and select **Storage > New Storage Repository**.
6. Create a new SR with the appropriate parameters required to reconnect the existing SR to the destination host. In the case of moving a SR between pools or hosts within a site the parameters may be identical to the source pool.

7. Every time a new SR is created the storage is checked to see if it contains an existing SR. If so, an option is presented allowing re-attachment of the existing SR. If this option is not displayed the parameters specified during SR creation are not correct.
8. Select **Reattach**.
9. Select the new SR in the tree view and then select the **Storage** tab to view the existing VDIs present on the SR.
10. In xsconsole on the destination host, select the **Backup, Restore, and Update** menu option, select the **Restore Virtual Machine Metadata** option, and select the newly re-attached SR.
11. The VDIs on the selected SR are inspected to find the metadata VDI. Once found, select the metadata backup you want to use.
12. Select the **Only VMs on this SR** option to restore the VMs.

Note:

Use the **All VM Metadata** option when moving multiple SRs between hosts or pools, or when using tiered storage where VMs to be restored have VDIs on multiple SRs. When using this option ensure all required SRs have been reattached to the destination host prior running the restore.

13. The VMs are restored in the destination pool in a shutdown state and are available for use.

Using Portable SRs for Manual Multi-Site Disaster Recovery

The Portable SR feature can be used in combination with storage layer replication in order to simplify the process of creating and enabling a disaster recovery (DR) site. Using storage layer replication to mirror or replicate LUNs that comprise portable SRs between production and DR sites allows all required data to be automatically present in the DR site. The constraints that apply when moving portable SRs between hosts or pools within the same site also apply in the multi-site case, but the production and DR sites are not required to have the same number of hosts. This allows use of either dedicated DR facilities or non-dedicated DR sites that run other production workloads.

Using portable SRs with storage layer replication between sites to enable the DR site in case of disaster

1. Any storage layer configuration required to enable the mirror or replica LUN in the DR site are performed.
2. An SR is created for each LUN in the DR site.
3. VMs are restored from metadata on one or more SRs.
4. Any adjustments to VM configuration required by differences in the DR site, such as IP addressing, are performed.
5. VMs are started and verified.
6. Traffic is routed to the VMs in the DR site.

VM Snapshots

XenServer provides a convenient snapshotting mechanism that can take a snapshot of a VM storage and metadata at a given time. Where necessary, IO is temporarily halted while the snapshot is being taken to ensure that a self-consistent disk image can be captured.

Snapshot operations result in a snapshot VM that is similar to a template. The VM snapshot contains all the storage information and VM configuration, including attached VIFs, allowing them to be exported and restored for backup purposes. Snapshots are supported on all storage types, though for the LVM-based



storage types the storage repository must have been upgraded if it was created on a previous version of XenServer and the volume must be in the default format (`type=raw` volumes cannot be snapshotted).

The snapshotting operation is a 2 step process:

- Capturing metadata as a template.
- Creating a VDI snapshot of the disk(s).

Three types of VM snapshots are supported: regular, quiesced, and snapshot with memory

Regular Snapshots

Regular snapshots are crash consistent and can be performed on all VM types, including Linux VMs.

Quiesced Snapshots

Quiesced snapshots take advantage of the Windows Volume Shadow Copy Service (VSS) to generate application consistent point-in-time snapshots. The VSS framework helps VSS-aware applications (for example Microsoft Exchange or Microsoft SQL Server) flush data to disk and prepare for the snapshot before it is taken.

Quiesced snapshots are therefore safer to restore, but can have a greater performance impact on a system while they are being taken. They may also fail under load so more than one attempt to take the snapshot may be required.

XenServer supports quiesced snapshots on Windows Server 2003 and Windows Server 2008 for both 32-bit and 64-bit variants. Windows 2000, Windows XP and Windows Vista are not supported. For further detail on quiesced snapshots, refer to the Advanced Notes for Quiesced Snapshots section later in this chapter.

Snapshots with memory

In addition to saving the VMs memory (storage) and metadata, snapshots with memory also save the VMs state (RAM). This can be useful if you are upgrading or patching software, or want to test a new application, but also want the option to be able to get back to the current, pre-change state (RAM) of the VM. Reverting back to a snapshot with memory, does not require a reboot of the VM.

You can take a snapshot with memory of a running or suspended VM via the XenAPI, the `xe` CLI, or by using XenCenter.

Creating a VM Snapshot

Before taking a snapshot, see the section called “Preparing to clone a Windows VM” in XenServer Virtual Machine Installation Guide and the section called “Preparing to clone a Linux VM” in XenServer Virtual Machine Installation Guide for information about any special operating system-specific configuration and considerations to take into account.

Firstly, ensure that the VM is running or suspended so that the memory status can be captured. The simplest way to select the VM on which the operation is to be performed is by supplying the argument `vm=<name>` or `vm=<vm uuid>` .

Run the **vm-snapshot** and **vm-snapshot-with-quiesce** commands to take a snapshot of a VM.

```
xe vm-snapshot vm=<vm uuid> new-name-label=<vm_snapshot_name>
xe vm-snapshot-with-quiesce vm=<vm uuid> new-name-label=<vm_snapshot_name>
```



Creating a snapshot with memory

Run the **vm-checkpoint** command, giving a descriptive name for the snapshot with memory, so that you can identify it later:

```
xe vm-checkpoint vm=<vm uuid> new-name-label=<name of the checkpoint>
```

When XenServer has completed creating the snapshot with memory - its uuid will be displayed.

For example:

```
xe vm-checkpoint vm=2d1d9a08-e479-2f0a-69e7-24a0e062dd35 \
  new-name-label=example_checkpoint_1
b3c0f369-59a1-dd16-ecd4-a1211df29886
```

A snapshot with memory requires at least 4MB of disk space per disk, plus the size of the RAM, plus around 20% overhead. So a checkpoint with 256MB RAM would require approximately 300MB of storage.

Note:

During the checkpoint creation process, the VM is paused for a brief period of time, and cannot be used during this period.

To list all of the snapshots on a XenServer pool

Run the **snapshot-list** command:

```
xe snapshot-list
```

This lists all of the snapshots in the XenServer pool.

To list the snapshots on a particular VM

You will need to know the uuid of the particular VM; to do this run the **vm-list** command.

```
xe vm-list
```

This displays a list of all VMs and their UUIDs. For example:

```
xe vm-list
uuid ( RO): 116dd310-a0ef-a830-37c8-df41521ff72d
name-label ( RW): Windows Server 2003 (1)
power-state ( RO): halted

uuid ( RO): 96fde888-2a18-c042-491a-014e22b07839
name-label ( RW): Windows XP SP3 (1)
power-state ( RO): running

uuid ( RO): dff45c56-426a-4450-a094-d3bba0a2ba3f
name-label ( RW): Control domain on host
power-state ( RO): running
```

VMs can also be specified by filtering the full list of VMs on the values of fields.



For example, specifying **power-state=halted** will select all VMs whose power-state field is equal to 'halted'. Where multiple VMs are matching, the option **--multiple** must be specified to perform the operation. The full list of fields that can be matched can be obtained by the command **xe vm-list params=all**.

Locate the required VM and then enter the following:

```
xe snapshot-list snapshot-of=<vm uuid>
```

For example:

```
xe snapshot-list snapshot-of=2d1d9a08-e479-2f0a-69e7-24a0e062dd35
```

This lists the snapshots currently on that VM:

```
uuid ( RO): d7eefb03-39bc-80f8-8d73-2ca1bab7dcff
name-label ( RW): Regular
name-description ( RW):
snapshot_of ( RO): 2d1d9a08-e479-2f0a-69e7-24a0e062dd35
snapshot_time ( RO): 20090914T15:37:00Z

uuid ( RO): 1760561d-a5d1-5d5e-2be5-d0dd99a3b1ef
name-label ( RW): Snapshot with memory
name-description ( RW):
snapshot_of ( RO): 2d1d9a08-e479-2f0a-69e7-24a0e062dd35
snapshot_time ( RO): 20090914T15:39:45Z
```

Restoring a VM to its previous state

Ensure that you have the uuid of the snapshot that you want to revert to, and then run the **snapshot-revert** command:

To do this:

1. Run the **snapshot-list** command to find the UUID of the snapshot or checkpoint that you want to revert to:

```
xe snapshot-list
```

2. Note the uuid of the snapshot, and then run the following command to revert:

```
xe snapshot-revert snapshot-uuid=<snapshot uuid>
```

For example:

```
xe snapshot-revert snapshot-uuid=b3c0f369-59a1-dd16-ecd4-a1211df29886
```

After reverting to a checkpoint, the VM will be suspended.

Note:

It is possible to revert to any snapshot in time, forwards or backwards. Existing snapshots and checkpoints are not deleted during revert.

Deleting a snapshot

Ensure that you have the UUID of the checkpoint or snapshot that you wish to remove, and then run the following command:



1. Run the **snapshot-list** command to find the UUID of the snapshot or checkpoint that you want to revert to:

```
xe snapshot-list
```

2. Note the UUID of the snapshot, and then run the **snapshot-uninstall** command to remove it:

```
xe snapshot-uninstall snapshot-uuid=<snapshot-uuid>
```

3. This command alerts you to the VM and VDI's that will be deleted. Type **yes** to confirm.

For example:

```
xe snapshot-uninstall snapshot-uuid=1760561d-a5d1-5d5e-2be5-d0dd99a3b1ef
The following items are about to be destroyed
VM : 1760561d-a5d1-5d5e-2be5-d0dd99a3b1ef (Snapshot with memory)
VDI: 11a4aa81-3c6b-4f7d-805a-b6ea02947582 (0)
VDI: 43c33fe7-a768-4612-bf8c-c385e2c657ed (1)
VDI: 4c33c84a-a874-42db-85b5-5e29174fa9b2 (Suspend image)
Type 'yes' to continue
yes
All objects destroyed
```

If you only want to remove the metadata of a checkpoint or snapshot, run the following command:

```
xe snapshot-destroy snapshot-uuid=<snapshot-uuid>
```

For example:

```
xe snapshot-destroy snapshot-uuid=d7eefb03-39bc-80f8-8d73-2ca1bab7dcff
```

Snapshot Templates

Creating a template from a snapshot

You can create a VM template from a snapshot, however its memory state will be removed.

To do this:

1. Use the command **snapshot-copy** and specify a **new-name-label** for the template:

```
xe snapshot-copy new-name-label=<vm-template-name> \
  snapshot-uuid=<uuid of the snapshot>
```

For example:

```
xe snapshot-copy new-name-label=example_template_1
  snapshot-uuid=b3c0f369-59a1-dd16-ecd4-a1211df29886
```

Note:

This creates a template object in the SAME pool. This template exists in the XenServer database for the current pool only.

2. To verify that the template has been created, run the command **template-list**:

```
xe template-list
```



This will list all of the templates on the XenServer host.

Exporting a snapshot to a template

When you export a VM snapshot, a complete copy of the VM (including disk images) is stored as a single file on your local machine, with a **.xva** file extension.

To do this:

1. Use the command **snapshot-export-to-template** to create a new template file:

```
xe snapshot-export-to-template snapshot-uuid=<snapshot-uuid> \  
    filename=<template- filename>
```

For example:

```
xe snapshot-export-to-template snapshot-uuid=b3c0f369-59a1-dd16-ecd4-a1211df29886 \  
    filename=example_template_export
```

The VM export/import feature can be used in a number of different ways:

- As a convenient backup facility for your VMs. An exported VM file can be used to recover an entire VM in the event of disaster.
- As a way of quickly copying a VM, for example, a special-purpose server configuration that you use many times. You simply configure the VM the way you want it, export it, and then import it to create copies of your original VM.
- As a simple method for moving a VM to another server.

For further information on the use of templates refer to the Creating VMs chapter in the XenServer Virtual Machine Installation Guide and also the Managing virtual machines section in the XenCenter Help.

Advanced Notes for Quiesced Snapshots

Note:

Do not forget to install the Xen VSS provider in the Windows guest in order to support VSS. This is done using the `install-XenProvider.cmd` script provided with the Windows PV drivers. More details can be found in the Virtual Machine Installation Guide in the Windows section.

In general, a VM can only access VDI snapshots (not VDI clones) of itself using the VSS interface. There is a flag that can be set by the XenServer administrator whereby adding an attribute of `snapmanager=true` to the VM's `other-config` allows that VM to import snapshots of VDIs from other VMs.

Warning:

This opens a security vulnerability and should be used with care. This feature allows an administrator to attach VSS snapshots using an in-guest transportable snapshot ID as generated by the VSS layer to another VM for the purposes of backup.

VSS quiesce timeout: the Microsoft VSS quiesce period is set to a non-configurable value of 10 seconds, and it is quite probable that a snapshot may not be able to complete in time. If, for example the XAPI daemon has queued additional blocking tasks such as an SR scan, the VSS snapshot may timeout and fail. The operation should be retried if this happens.

Note:

The more VBDs attached to a VM, the more likely it is that this timeout may be reached. Citrix recommends attaching no more than 2 VBDs to a VM to avoid reaching the timeout. However, there is a workaround to this problem. The probability of taking a successful VSS based snapshot of a VM with more than 2 VBDs can be increased manifold, if all the VDIs for the VM are hosted on different SRs.

VSS snapshot all the disks attached to a VM: in order to store all data available at the time of a VSS snapshot, the XAPI manager will snapshot all disks and the VM metadata associated with a VM that can be snapshotted using the XenServer storage manager API. If the VSS layer requests a snapshot of only a subset of the disks, a full VM snapshot will not be taken.

vm-snapshot-with-quiesce produces bootable snapshot VM images: To achieve this, the XenServer VSS hardware provider makes snapshot volumes writable, including the snapshot of the boot volume.

VSS snap of volumes hosted on dynamic disks in the Windows Guest : The **vm-snapshot-with-quiesce CLI** and the XenServer VSS hardware provider do not support snapshots of volumes hosted on dynamic disks on the Windows VM.

Note:

Using EqualLogic or NetApp storage requires a **Citrix XenServer Advanced Edition or higher** license. To learn more about XenServer editions, and to find out how to upgrade, visit the Citrix website [here](#).

Note:

Do not forget to install the Xen VSS provider in the Windows guest in order to support VSS. This is done using the `install-XenProvider.cmd` script provided with the Windows PV drivers. More details can be found in the *Virtual Machine Installation Guide* in the Windows section.

In general, a VM can only access VDI snapshots (not VDI clones) of itself using the VSS interface. There is a flag that can be set by the XenServer administrator whereby adding an attribute of `snapmanager=true` to the VM's `other-config` allows that VM to import snapshots of VDIs from other VMs.

Warning:

This opens a security vulnerability and should be used with care. This feature allows an administrator to attach VSS snapshots using an in-guest transportable snapshot ID as generated by the VSS layer to another VM for the purposes of backup.

VSS quiesce timeout: the Microsoft VSS quiesce period is set to a non-configurable value of 10 seconds, and it is quite probable that a snapshot may not be able to complete in time. If, for example the XAPI daemon has queued additional blocking tasks such as an SR scan, the VSS snapshot may timeout and fail. The operation should be retried if this happens.

Note:

The more VBDs attached to a VM, the more likely it is that this timeout may be reached. Citrix recommends attaching no more than 2 VBDs to a VM to avoid reaching the timeout. However, there is a workaround to this problem. The probability of taking a successful VSS based snapshot of a VM with more than 2 VBDs can be increased manifold, if all the VDIs for the VM are hosted on different SRs.

VSS snapshot all the disks attached to a VM: in order to store all data available at the time of a VSS snapshot, the XAPI manager will snapshot all disks and the VM metadata associated with a VM that can be snapshotted using the XenServer storage manager API. If the VSS layer requests a snapshot of only a subset of the disks, a full VM snapshot will not be taken.

vm-snapshot-with-quiesce produces bootable snapshot VM images : To achieve this, the XenServer VSS hardware provider makes snapshot volumes writable, including the snapshot of the boot volume.

VSS snap of volumes hosted on dynamic disks in the Windows Guest : The `vm-snapshot-with-quiesce` CLI and the XenServer VSS hardware provider do not support snapshots of volumes hosted on dynamic disks on the Windows VM.

Coping with machine failures

This section provides details of how to recover from various failure scenarios. All failure recovery scenarios require the use of one or more of the backup types listed in [the section called “Backups”](#).

Member failures

In the absence of HA, master nodes detect the failures of members by receiving regular heartbeat messages. If no heartbeat has been received for 200 seconds, the master assumes the member is dead. There are two ways to recover from this problem:

- Repair the dead host (e.g., by physically rebooting it). When the connection to the member is restored, the master will mark the member as alive again.
- Shutdown the host and instruct the master to forget about the member node using the **`xe host-forget`** CLI command. Once the member has been forgotten, all the VMs which were running there will be marked as offline and can be restarted on other XenServer hosts. Note it is very important to ensure that the XenServer host is actually offline, otherwise VM data corruption might occur. Be careful not to split your pool into multiple pools of a single host by using **`xe host-forget`**, since this could result in them all mapping the same shared storage and corrupting VM data.

Warning:

- If you are going to use the forgotten host as a XenServer host again, perform a fresh installation of the XenServer software.
- Do not use **`xe host-forget`** command if HA is enabled on the pool. Disable HA first, then forget the host, and then re-enable HA.

When a member XenServer host fails, there may be VMs still registered in the *running* state. If you are sure that the member XenServer host is definitely down, and that the VMs have not been brought up on another XenServer host in the pool, use the **`xe vm-reset-powerstate`** CLI command to set the power state of the VMs to `halted`. See [the section called “vm-reset-powerstate”](#) for more details.

Warning:

Incorrect use of this command can lead to data corruption. Only use this command if absolutely necessary.

Master failures

Every member of a resource pool contains all the information necessary to take over the role of master if required. When a master node fails, the following sequence of events occurs:

1. If HA is enabled, another master is elected automatically.
2. If HA is not enabled, each member will wait for the master to return.

If the master comes back up at this point, it re-establishes communication with its members, and operation returns to normal.

If the master is really dead, choose one of the members and run the command **`xe pool-emergency-transition-to-master`** on it. Once it has become the master, run the command **`xe pool-recover-slaves`** and the members will now point to the new master.



If you repair or replace the server that was the original master, you can simply bring it up, install the XenServer host software, and add it to the pool. Since the XenServer hosts in the pool are enforced to be homogeneous, there is no real need to make the replaced server the master.

When a member XenServer host is transitioned to being a master, you should also check that the default pool storage repository is set to an appropriate value. This can be done using the **xe pool-param-list** command and verifying that the *default-SR* parameter is pointing to a valid storage repository.

Pool failures

In the unfortunate event that your entire resource pool fails, you will need to recreate the pool database from scratch. Be sure to regularly back up your pool-metadata using the **xe pool-dump-database** CLI command (see [the section called “pool-dump-database”](#)).

To restore a completely failed pool

1. Install a fresh set of hosts. Do not pool them up at this stage.
2. For the host nominated as the master, restore the pool database from your backup using the **xe pool-restore-database** (see [the section called “pool-restore-database”](#)) command.
3. Connect to the master host using XenCenter and ensure that all your shared storage and VMs are available again.
4. Perform a pool join operation on the remaining freshly installed member hosts, and start up your VMs on the appropriate hosts.

Coping with Failure due to Configuration Errors

If the physical host machine is operational but the software or host configuration is corrupted:

To restore host software and configuration

1. Run the command:

```
xe host-restore host=<host> file-name=<hostbackup>
```

2. Reboot to the host installation CD and select **Restore from backup**.

Physical Machine failure

If the physical host machine has failed, use the appropriate procedure listed below to recover.

Warning:

Any VMs which were running on a previous member (or the previous host) which has failed will still be marked as `Running` in the database. This is for safety-- simultaneously starting a VM on two different hosts would lead to severe disk corruption. If you are sure that the machines (and VMs) are offline you can reset the VM power state to `Halted`:

```
xe vm-reset-powerstate vm=<vm_uuid> --force
```

VMs can then be restarted using XenCenter or the CLI.

Replacing a failed master with a still running member

1. Run the commands:

```
xe pool-emergency-transition-to-master
xe pool-recover-slaves
```

2. If the commands succeed, restart the VMs.

To restore a pool with all hosts failed

1. Run the command:

```
xe pool-restore-database file-name=<backup>
```

Warning:

This command will only succeed if the target machine has an appropriate number of appropriately named NICs.

2. If the target machine has a different view of the storage (for example, a block-mirror with a different IP address) than the original machine, modify the storage configuration using the **pbd-destroy** command and then the **pbd-create** command to recreate storage configurations. See [the section called “PBD commands”](#) for documentation of these commands.
3. If you have created a new storage configuration, use **pbd-plug** or **Storage > Repair Storage Repository** menu item in XenCenter to use the new configuration.
4. Restart all VMs.

To restore a VM when VM storage is not available

1. Run the command:

```
xe vm-import filename=<backup> --metadata
```

2. If the metadata import fails, run the command:

```
xe vm-import filename=<backup> --metadata --force
```

This command will attempt to restore the VM metadata on a 'best effort' basis.

3. Restart all VMs.

Monitoring and managing XenServer

XenServer and XenCenter provide access to alerts that are generated when noteworthy things happen. XenCenter provides various mechanisms of grouping and maintaining metadata about managed VMs, hosts, storage repositories, and so on.

Note:

Full monitoring and alerting functionality is only available with XenServer Advanced edition or above. To find out about XenServer editions and how to upgrade, visit the Citrix website [here](#).

Alerts

XenServer generates alerts for the following events.

Configurable Alerts:

- New XenServer patches available
- New XenServer version available
- New XenCenter version available

Alerts generated by XenCenter:

Alert	Description
XenCenter old	the XenServer expects a newer version but can still connect to the current version
XenCenter out of date	XenCenter is too old to connect to XenServer
XenServer out of date	XenServer is an old version that the current XenCenter cannot connect to
License expired alert	your XenServer license has expired
Missing IQN alert	XenServer uses iSCSI storage but the host IQN is blank
Duplicate IQN alert	XenServer uses iSCSI storage, and there are duplicate host IQNs

Alerts generated by XenServer:

- ha_host_failed
- ha_host_was_fenced
- ha_network_bonding_error
- ha_pool_drop_in_plan_exists_for
- ha_pool_overcommitted
- ha_protected_vm_restart_failed
- ha_statefile_lost
- host_clock_skew_detected
- host_sync_data_failed



- license_does_not_support_pooling
- pbd_plug_failed_on_server_start
- pool_master_transition

The following alerts appear on the performance graphs in XenCenter. See the *XenCenter online help* for more information:

- vm_cloned
- vm_crashed
- vm_rebooted
- vm_resumed
- vm_shutdown
- vm_started
- vm_suspended

Customizing Alerts

The performance monitoring `perfmon` runs once every 5 minutes and requests updates from XenServer which are averages over 1 minute, but these defaults can be changed in `/etc/sysconfig/perfmon`.

Every 5 minutes `perfmon` reads updates of performance variables exported by the XAPI instance running on the same host. These variables are separated into one group relating to the host itself, and a group for each VM running on that host. For each VM and also for the host, `perfmon` reads in the `other-config:perfmon` parameter and uses this string to determine which variables it should monitor, and under which circumstances to generate a message.

`vm:other-config:perfmon` and `host:other-config:perfmon` values consist of an XML string like the one below:

```
<config>
  <variable>
    <name value="cpu_usage"/>
    <alarm_trigger_level value="LEVEL"/>
  </variable>
  <variable>
    <name value="network_usage"/>
    <alarm_trigger_level value="LEVEL"/>
  </variable>
</config>
```

Valid VM Elements

name

what to call the variable (no default). If the name value is one of `cpu_usage`, `network_usage`, or `disk_usage` the `rrd_regex` and `alarm_trigger_sense` parameters are not required as defaults for these values will be used.

alarm_priority

the priority of the messages generated (default 5)

alarm_trigger_level

level of value that triggers an alarm (no default)



alarm_trigger_sense

high if *alarm_trigger_level* is a maximum value otherwise low if the *alarm_trigger_level* is a minimum value. (default high)

alarm_trigger_period

number of seconds that values above or below the alarm threshold can be received before an alarm is sent (default 60)

alarm_auto_inhibit_period

number of seconds this alarm disabled after an alarm is sent (default 3600)

consolidation_fn

how to combine variables from *rrd_updates* into one value (default is *sum* - other choice is *average*)

rrd_regex

regular expression to match the names of variables returned by the **xe vm-data-source-list uuid=<vmuuid>** command that should be used to compute the statistical value. This parameter has defaults for the named variables *cpu_usage*, *network_usage*, and *disk_usage*. If specified, the values of all items returned by **xe vm-data-source-list** whose names match the specified regular expression will be consolidated using the method specified as the *consolidation_fn*.

Valid Host Elements

name

what to call the variable (no default)

alarm_priority

the priority of the messages generated (default 5)

alarm_trigger_level

level of value that triggers an alarm (no default)

alarm_trigger_sense

high if *alarm_trigger_level* is a maximum value otherwise low if the *alarm_trigger_level* is a minimum value. (default high)

alarm_trigger_period

number of seconds that values above or below the alarm threshold can be received before an alarm is sent (default 60)

alarm_auto_inhibit_period

number of seconds this alarm disabled after an alarm is sent (default 3600)

consolidation_fn

how to combine variables from **rrd_updates** into one value (default *sum* - other choice is *average*)

rrd_regex

regular expression to match the names of variables returned by the **xe vm-data-source-list uuid=<vmuuid>** command that should be used to compute the statistical value. This parameter has defaults for the named variables *cpu_usage* and *network_usage*. If specified, the values of all items returned by **xe vm-data-source-list** whose names match the specified regular expression will be consolidated using the method specified as the *consolidation_fn*.

Configuring Email Alerts

Note:

Email alerts are only available in a pool with a is only available with XenServer Advanced edition or above. To find out about XenServer editions and how to upgrade, visit the Citrix website [here](#).



Alerts generated from XenServer can also be automatically e-mailed to the resource pool administrator, in addition to being visible from the XenCenter GUI. To configure this, specify the email address and SMTP server:

```
pool:other-config:mail-destination=<joe.bloggs@domain.tld>  
pool:other-config:ssmtp-mailhub=<smtp.domain.tld[:port]>
```

You can also specify the minimum value of the priority field in the message before the email will be sent:

```
pool:other-config:mail-min-priority=<level>
```

The default priority level is 5.

Note:

Some SMTP servers only forward mails with addresses that use FQDNs. If you find that emails are not being forwarded it may be for this reason, in which case you can set the server hostname to the FQDN so this is used when connecting to your mail server.

Custom Fields and Tags

XenCenter supports the creation of tags and custom fields, which allows for organization and quick searching of VMs, storage and so on. See the XenCenter online help for more information.

Custom Searches

XenCenter supports the creation of customized searches. Searches can be exported and imported, and the results of a search can be displayed in the navigation pane. See the XenCenter online help for more information.

Determining throughput of physical bus adapters

For FC, SAS and iSCSI HBAs you can determine the network throughput of your PBDs using the following procedure.

To determine PBD throughput

1. List the PBDs on a host.
2. Determine which LUNs are routed over which PBDs.
3. For each PBD and SR, list the VBDs that reference VDIs on the SR.
4. For all active VBDs that are attached to VMs on the host, calculate the combined throughput.

For iSCSI and NFS storage, check your network statistics to determine if there is a throughput bottleneck at the array, or whether the PBD is saturated.

Troubleshooting

If you experience odd behavior, application crashes, or have other issues with a XenServer host, this chapter is meant to help you solve the problem if possible and, failing that, describes where the application logs are located and other information that can help your Citrix Solution Provider and Citrix track and resolve the issue.

Troubleshooting of installation issues is covered in the *XenServer Installation Guide*. Troubleshooting of Virtual Machine issues is covered in the *XenServer Virtual Machine Installation Guide*.

Important:

We recommend that you follow the troubleshooting information in this chapter solely under the guidance of your Citrix Solution Provider or Citrix Support.

Citrix provides two forms of support: you can receive free self-help support on the [Support site](#), or you may [purchase our Support Services](#) and directly submit requests by filing an online Support Case. Our free web-based resources include product documentation, a Knowledge Base, and discussion forums.

XenServer host logs

XenCenter can be used to gather XenServer host information. Click on **Get Server Status Report...** in the **Tools** menu to open the **Server Status Report** wizard. You can select from a list of different types of information (various logs, crash dumps, etc.). The information is compiled and downloaded to the machine that XenCenter is running on. For details, see the XenCenter Help.

Additionally, the XenServer host has several CLI commands to make it simple to collate the output of logs and various other bits of system information using the utility **xen-bugtool**. Use the **xe** command **host-bugreport-upload** to collect the appropriate log files and system information and upload them to the Citrix Support ftp site. Please refer to [the section called “host-bugreport-upload”](#) for a full description of this command and its optional parameters. If you are requested to send a crashdump to Citrix Support, use the **xe** command **host-crashdump-upload**. Please refer to [the section called “host-crashdump-upload”](#) for a full description of this command and its optional parameters.

It is possible that sensitive information might be written into the XenServer host logs.

By default, the server logs report only errors and warnings. If you need to see more detailed information, you can enable more verbose logging. To do so, use the **host-loglevel-set** command:

host-loglevel-set *log-level=level*

where *level* can be 0, 1, 2, 3, or 4, where 0 is the most verbose and 4 is the least verbose.

Log files greater than 5 MB are rotated, keeping 4 revisions. The **logrotate** command is run hourly.

Sending host log messages to a central server

Rather than have logs written to the control domain filesystem, you can configure a XenServer host to write them to a remote server. The remote server must have the syslogd daemon running on it to receive the logs and aggregate them correctly. The syslogd daemon is a standard part of all flavors of Linux and Unix, and third-party versions are available for Windows and other operating systems.

To write logs to a remote server

1. Set the **syslog_destination** parameter to the hostname or IP address of the remote server where you want the logs to be written:



```
xe host-param-set uuid=<xenserver_host_uuid> logging:syslog_destination=<hostname>
```

2. Issue the command:

```
xe host-syslog-reconfigure uuid=<xenserver_host_uuid>
```

to enforce the change. (You can also execute this command remotely by specifying the *host* parameter.)

XenCenter logs

XenCenter also has client-side log. This file includes a complete description of all operations and errors that occur when using XenCenter. It also contains informational logging of events that provide you with an audit trail of various actions that have occurred. The XenCenter log file is stored in your profile folder. If XenCenter is installed on Windows XP, the path is

```
%userprofile%\AppData\Citrix\XenCenter\logs\XenCenter.log
```

If XenCenter is installed on Windows Vista, the path is

```
%userprofile%\AppData\Citrix\Roaming\XenCenter\logs\XenCenter.log
```

To quickly locate the XenCenter log files, for example, when you want to open or email the log file, click on **View Application Log Files** in the XenCenter **Help** menu.

Troubleshooting connections between XenCenter and the XenServer host

If you have trouble connecting to the XenServer host with XenCenter, check the following:

- Is your XenCenter an older version than the XenServer host you are attempting to connect to?

The XenCenter application is backward-compatible and can communicate properly with older XenServer hosts, but an older XenCenter cannot communicate properly with newer XenServer hosts.

To correct this issue, install a XenCenter version that is the same, or newer, than the XenServer host version.

- Is your license current?

You can see the expiration date for your License Key in the XenServer host **General** tab under the **Licenses** section in XenCenter.

Also, if you upgraded your software from version 3.2.0 to the current version, you should also have received and applied a new License file.

For details on licensing a host, see the chapter "XenServer Licensing" in the *XenServer Installation Guide*.

- The XenServer host talks to XenCenter using HTTPS over port 443 (a two-way connection for commands and responses using the XenAPI), and 5900 for graphical VNC connections with paravirtual Linux VMs. If you have a firewall enabled between the XenServer host and the machine running the client software, make sure that it allows traffic from these ports.

Appendix A. Command line interface

This chapter describes the XenServer command line interface (CLI). The `xe` CLI enables the writing of scripts for automating system administration tasks and allows integration of XenServer into an existing IT infrastructure.

The `xe` command line interface is installed by default on XenServer hosts and is included with XenCenter. A stand-alone remote CLI is also available for Linux.

On Windows, the `xe.exe` CLI executable is installed along with XenCenter.

To use it, open a Windows Command Prompt and change directories to the directory where the file resides (typically `C:\Program Files\XenSource\XenCenter`), or add its installation location to your system path.

On Linux, you can install the stand-alone `xe` CLI executable from the RPM named `xe-cli-5.6-@BUILD_NUMBER@.i386.rpm` on the Linux Pack CD, as follows:

```
rpm -ivh xe-cli-5.6-@BUILD_NUMBER@.i386.rpm
```

Basic help is available for CLI commands on-host by typing:

```
xe help command
```

A list of the most commonly-used `xe` commands is displayed if you type:

```
xe help
```

or a list of all `xe` commands is displayed if you type:

```
xe help --all
```

Basic `xe` syntax

The basic syntax of all XenServer `xe` CLI commands is:

```
xe <command-name> <argument=value> <argument=value> ...
```

Each specific command contains its own set of arguments that are of the form `argument=value`. Some commands have required arguments, and most have some set of optional arguments. Typically a command will assume default values for some of the optional arguments when invoked without them.

If the `xe` command is executed remotely, additional connection and authentication arguments are used. These arguments also take the form `argument=argument_value`.

The `server` argument is used to specify the hostname or IP address. The `username` and `password` arguments are used to specify credentials. A `password-file` argument can be specified instead of the password directly. In this case an attempt is made to read the password from the specified file (stripping CRs and LFs off the end of the file if necessary), and use that to connect. This is more secure than specifying the password directly at the command line.

The optional `port` argument can be used to specify the agent port on the remote XenServer host (defaults to 443).



Example: On the local XenServer host:

```
xe vm-list
```

Example: On the remote XenServer host:

```
xe vm-list -user <username> -password <password> -server <hostname>
```

Shorthand syntax is also available for remote connection arguments:

-u	username
-pw	password
-pwf	password file
-p	port
-s	server

Example: On a remote XenServer host:

```
xe vm-list -u <myuser> -pw <mypassword> -s <hostname>
```

Arguments are also taken from the environment variable `XE_EXTRA_ARGS`, in the form of comma-separated key/value pairs. For example, in order to enter commands on one XenServer host that are run on a remote XenServer host, you could do the following:

```
export XE_EXTRA_ARGS="server=jeffbeck,port=443,username=root,password=pass"
```

This command means that you will not need to specify the remote XenServer host parameters anymore, in each `xe` command you execute.

Using the `XE_EXTRA_ARGS` environment variable also enables tab completion of `xe` commands when issued against a remote XenServer host, which is disabled by default.

Special characters and syntax

To specify argument/value pairs on the **xe** command line, write:

argument=value

Unless the value includes spaces, do not use quotes. There should be no whitespace in between the argument name, the equals sign (=), and the value. Any argument not conforming to this format will be ignored.

For values containing spaces, write:

argument="value with spaces"

If you use the CLI while logged into a XenServer host, commands have a tab completion feature similar to that in the standard Linux bash shell. If you type, for example **xe vm-l** and then press the **TAB** key, the rest of the command will be displayed when it is unambiguous. If more than one command begins with **vm-**



I, pressing **TAB** a second time will list the possibilities. This is particularly useful when specifying object UUIDs in commands.

Note:

When executing commands on a remote XenServer host, tab completion does not normally work. However if you put the server, username, and password in an environment variable called `XE_EXTRA_ARGS` on the machine from which you are entering the commands, tab completion is enabled. See [the section called “Basic xe syntax”](#) for details.

Command types

Broadly speaking, the CLI commands can be split in two halves: Low-level commands concerned with listing and parameter manipulation of API objects, and higher level commands for interacting with VMs or hosts in a more abstract level. The low-level commands are:

- `<class>-list`
- `<class>-param-get`
- `<class>-param-set`
- `<class>-param-list`
- `<class>-param-add`
- `<class>-param-remove`
- `<class>-param-clear`

where `<class>` is one of:

- bond
- console
- host
- host-crashdump
- host-cpu
- network
- patch
- pbd
- pif
- pool
- sm
- sr
- task
- template
- vbd
- vdi
- vif
- vlan

- vm

Note that not every value of `<class>` has the full set of `<class>-param-` commands; some have just a subset.

Parameter types

The objects that are addressed with the `xe` commands have sets of parameters that identify them and define their states.

Most parameters take a single value. For example, the `name-label` parameter of a VM contains a single string value. In the output from parameter list commands such as `xe vm-param-list`, such parameters have an indication in parentheses that defines whether they can be read and written to, or are read-only. For example, the output of `xe vm-param-list` on a specified VM might have the lines

```
user-version ( RW): 1
is-control-domain ( RO): false
```

The first parameter, `user-version`, is writable and has the value 1. The second, `is-control-domain`, is read-only and has a value of false.

The two other types of parameters are multi-valued. A `set` parameter contains a list of values. A `map` parameter is a set of key/value pairs. As an example, look at the following excerpt of some sample output of the `xe vm-param-list` on a specified VM:

```
platform (MRW): acpi: true; apic: true; pae: true; nx: false
allowed-operations (SRO): pause; clean_shutdown; clean_reboot; \
hard_shutdown; hard_reboot; suspend
```

The `platform` parameter has a list of items that represent key/value pairs. The key names are followed by a colon character (:). Each key/value pair is separated from the next by a semicolon character (;). The M preceding the RW indicates that this is a map parameter and is readable and writable. The `allowed-operations` parameter has a list that makes up a set of items. The S preceding the RO indicates that this is a set parameter and is readable but not writable.

In `xe` commands where you want to filter on a map parameter, or set a map parameter, use the separator : (colon) between the map parameter name and the key/value pair. For example, to set the value of the `foo` key of the `other-config` parameter of a VM to `baa`, the command would be

```
xe vm-param-set uuid=<VM uuid> other-config:foo=baa
```

Note:

In previous releases the separator - (dash) was used in specifying map parameters. This syntax still works but is deprecated.

Low-level param commands

There are several commands for operating on parameters of objects: `<class>-param-get`, `<class>-param-set`, `<class>-param-add`, `<class>-param-remove`, `<class>-param-clear`, and `<class>-param-list`. Each of these takes a `uuid` parameter to specify the particular object. Since these are considered low-level commands, they must be addressed by UUID and not by the VM name label.

```
<class>-param-list uuid=<uuid>
```

Lists all of the parameters and their associated values. Unlike the `class-list` command, this will list the values of "expensive" fields.



`<class>-param-get uuid=<uuid> param-name=<parameter> [param-key=<key>]`

Returns the value of a particular parameter. If the parameter is a map, specifying the param-key will get the value associated with that key in the map. If param-key is not specified, or if the parameter is a set, it will return a string representation of the set or map.

`<class>-param-set uuid=<uuid> param=<value>...`

Sets the value of one or more parameters.

`<class>-param-add uuid=<uuid> param-name=<parameter> [<key>=<value>...] [param-key=<key>]`

Adds to either a map or a set parameter. If the parameter is a map, add key/value pairs using the `<key>=<value>` syntax. If the parameter is a set, add keys with the `<param-key>=<key>` syntax.

`<class>-param-remove uuid=<uuid> param-name=<parameter> param-key=<key>`

Removes either a key/value pair from a map, or a key from a set.

`<class>-param-clear uuid=<uuid> param-name=<parameter>`

Completely clears a set or a map.

Low-level list commands

The `<class>-list` command lists the objects of type `<class>`. By default it will list all objects, printing a subset of the parameters. This behavior can be modified in two ways: it can filter the objects so that it only outputs a subset, and the parameters that are printed can be modified.

To change the parameters that are printed, the argument *params* should be specified as a comma-separated list of the required parameters. For example:

```
xe vm-list params=name-label,other-config
```

Alternatively, to list all of the parameters, use the syntax:

```
xe vm-list params=all
```

Note that some parameters that are expensive to calculate will not be shown by the list command. These parameters will be shown as, for example:

```
allowed-VBD-devices (SRO): <expensive field>
```

To obtain these fields, use either the command `<class>-param-list` or `<class>-param-get`

To filter the list, the CLI will match parameter values with those specified on the command-line, only printing objects that match all of the specified constraints. For example:

```
xe vm-list HVM-boot-policy="BIOS order" power-state=halted
```

This command will only list those VMs for which *both* the field *power-state* has the value *halted*, and for which the field *HVM-boot-policy* has the value *BIOS order*.

It is also possible to filter the list based on the value of keys in maps, or on the existence of values in a set. The syntax for the first of these is **map-name:key=value**, and the second is **set-name:contains=value**



For scripting, a useful technique is passing `--minimal` on the command line, causing **xe** to print only the first field in a comma-separated list. For example, the command **xe vm-list --minimal** on a XenServer host with three VMs installed gives the three UUIDs of the VMs, for example:

```
a85d6717-7264-d00e-069b-3b1d19d56ad9,aaa3eec5-9499-bcf3-4c03-af10baea96b7, \
42c044de-df69-4b30-89d9-2c199564581d
```

xe command reference

This section provides a reference to the **xe** commands. They are grouped by objects that the commands address, and listed alphabetically.

Bonding commands

Commands for working with network bonds, for resilience with physical interface failover. See [the section called "Creating NIC bonds on a standalone host"](#) for details.

The bond object is a reference object which glues together *master* and *member* PIFs. The master PIF is the bonding interface which must be used as the overall PIF to refer to the bond. The member PIFs are a set of 2 or more physical interfaces which have been combined into the high-level bonded interface.

Bond parameters

Bonds have the following parameters:

Parameter Name	Description	Type
uuid	unique identifier/object reference for the bond	read only
master	UUID for the master bond PIF	read only
members	set of UUIDs for the underlying bonded PIFs	read only set parameter

bond-create

```
bond-create network-uuid=<network_uuid> pif-uuids=<pif_uuid_1,pif_uuid_2,...>
```

Create a bonded network interface on the network specified from a list of existing PIF objects. The command will fail if PIFs are in another bond already, if any member has a VLAN tag set, if the referenced PIFs are not on the same XenServer host, or if fewer than 2 PIFs are supplied.

bond-destroy

```
host-bond-destroy uuid=<bond_uuid>
```

Delete a bonded interface specified by its UUID from the XenServer host.

CD commands

Commands for working with physical CD/DVD drives on XenServer hosts.

CD parameters

CDs have the following parameters:

Parameter Name	Description	Type
uuid	unique identifier/object reference for the CD	read only
name-label	Name for the CD	read/write
name-description	Description text for the CD	read/write
allowed-operations	A list of the operations that can be performed on this CD	read only set parameter
current-operations	A list of the operations that are currently in progress on this CD	read only set parameter
sr-uuid	The unique identifier/object reference for the SR this CD is part of	read only
sr-name-label	The name for the SR this CD is part of	read only
vbd-uuids	A list of the unique identifiers for the VBDs on VMs that connect to this CD	read only set parameter
crashdump-uuids	Not used on CDs since crashdumps cannot be written to them	read only set parameter
virtual-size	Size of the CD as it appears to VMs (in bytes)	read only
physical-Utilization	amount of physical space that the CD image is currently taking up on the SR (in bytes)	read only
type	Set to <code>User</code> for CDs	read only
sharable	Whether or not the CD drive is sharable. Default is false.	read only
read-only	Whether the CD is read-only, if false, the device is writable. Always true for CDs.	read only
storage-lock	true if this disk is locked at the storage level	read only
parent	Reference to the parent disk, if this CD is part of a chain	read only

Parameter Name	Description	Type
missing	true if SR scan operation reported this CD as not present on disk	read only
other-config	A list of key/value pairs that specify additional configuration parameters for the CD	read/write map parameter
location	The path on which the device is mounted	read only
managed	true if the device is managed	read only
xenstore-data	Data to be inserted into the xenstore tree	read only map parameter
sm-config	names and descriptions of storage manager device config keys	read only map parameter
is-a-snapshot	True if this template is a CD snapshot	read only
snapshot_of	The UUID of the CD that this template is a snapshot of	read only
snapshots	The UUID(s) of any snapshots that have been taken of this CD	read only
snapshot_time	The timestamp of the snapshot operation	read only

cd-list

cd-list [params=<param1,param2,...>] [parameter=<parameter_value>...]

List the CDs and ISOs (CD image files) on the XenServer host or pool, filtering on the optional argument *params*.

If the optional argument *params* is used, the value of *params* is a string containing a list of parameters of this object that you want to display. Alternatively, you can use the keyword *all* to show all parameters. If *params* is not used, the returned list shows a default subset of all available parameters.

Optional arguments can be any number of the [CD parameters](#) listed at the beginning of this section.

Console commands

Commands for working with consoles.

The console objects can be listed with the standard object listing command (**xe console-list**), and the parameters manipulated with the standard parameter commands. See [the section called “Low-level param commands”](#) for details.

Console parameters

Consoles have the following parameters:

Parameter Name	Description	Type
uuid	The unique identifier/object reference for the console	read only
vm-uuid	The unique identifier/object reference of the VM this console is open on	read only
vm-name-label	The name of the VM this console is open on	read only
protocol	Protocol this console uses. Possible values are <i>vt100</i> : VT100 terminal, <i>rfb</i> : Remote FrameBuffer protocol (as used in VNC), or <i>rdp</i> : Remote Desktop Protocol	read only
location	URI for the console service	read only
other-config	A list of key/value pairs that specify additional configuration parameters for the console.	read/write map parameter

Event commands

Commands for working with events.

Event classes

Event classes are listed in the following table:

Class name	Description
pool	A pool of physical hosts
vm	A Virtual Machine
host	A physical host
network	A virtual network
vif	A virtual network interface
pif	A physical network interface (separate VLANs are represented as several PIFs)
sr	A storage repository
vdi	A virtual disk image
vbd	A virtual block device
pbd	The physical block devices through which hosts access SRs

event-wait

`event-wait class=<class_name> [<param-name>=<param_value>] [<param-name>!=<param_value>]`



Blocks other commands from executing until an object exists that satisfies the conditions given on the command line. `x=y` means "wait for field x to take value y", and `x!=y` means "wait for field x to take any value other than y".

Example: wait for a specific VM to be running

xe event-wait class=vm name-label=myvm power-state=running

Blocks other commands until a VM called `myvm` is in the `power-state` "running."

Example: wait for a specific VM to reboot:

xe event-wait class=vm uuid=\$VM start-time!= \$(xe vm-list uuid=\$VM params=start-time --minimal)

Blocks other commands until a VM with UUID `$VM` reboots (i.e. has a different `start-time` value).

The class name can be any of the [Event classes](#) listed at the beginning of this section, and the parameters can be any of those listed in the CLI command `class-param-list`.

Host (XenServer host) commands

Commands for interacting with XenServer host.

XenServer hosts are the physical servers running XenServer software. They have VMs running on them under the control of a special privileged Virtual Machine, known as the control domain or domain 0.

The XenServer host objects can be listed with the standard object listing command (**xe host-list**, **xe host-cpu-list**, and **xe host-crashdump-list**), and the parameters manipulated with the standard parameter commands. See [the section called "Low-level param commands"](#) for details.

Host selectors

Several of the commands listed here have a common mechanism for selecting one or more XenServer hosts on which to perform the operation. The simplest is by supplying the argument `host=<uuid_or_name_label>`. XenServer hosts can also be specified by filtering the full list of hosts on the values of fields. For example, specifying `enabled=true` will select all XenServer hosts whose `enabled` field is equal to `true`. Where multiple XenServer hosts are matching, and the operation can be performed on multiple XenServer hosts, the option `--multiple` must be specified to perform the operation. The full list of parameters that can be matched is described at the beginning of this section, and can be obtained by running the command **xe host-list params=all**. If no parameters to select XenServer hosts are given, the operation will be performed on all XenServer hosts.

Host parameters

XenServer hosts have the following parameters:

Parameter Name	Description	Type
uuid	The unique identifier/object reference for the XenServer host	read only
name-label	The name of the XenServer host	read/write
name-description	The description string of the XenServer host	read only

Parameter Name	Description	Type
enabled	false if disabled which prevents any new VMs from starting on them, which prepares the XenServer hosts to be shut down or rebooted; true if the host is currently enabled	read only
API-version-major	major version number	read only
API-version-minor	minor version number	read only
API-version-vendor	identification of API vendor	read only
API-version-vendor-implementation	details of vendor implementation	read only map parameter
logging	logging configuration	read/write map parameter
suspend-image-sr-uuid	the unique identifier/object reference for the SR where suspended images are put	read/write
crash-dump-sr-uuid	the unique identifier/object reference for the SR where crash dumps are put	read/write
software-version	list of versioning parameters and their values	read only map parameter
capabilities	list of Xen versions that the XenServer host can run	read only set parameter
other-config	A list of key/value pairs that specify additional configuration parameters for the XenServer host	read/write map parameter
hostname	XenServer host hostname	read only
address	XenServer host IP address	read only
supported-bootloaders	list of bootloaders that the XenServer host supports, for example, pygrub, eliloader	read only set parameter
memory-total	total amount of physical RAM on the XenServer host, in bytes	read only
memory-free	total amount of physical RAM remaining that can be allocated to VMs, in bytes	read only

Parameter Name	Description	Type
host-metrics-live	true if the host is operational	read only
logging	The <i>syslog_destination</i> key can be set to the hostname of a remote listening syslog service.	read/write map parameter
allowed-operations	lists the operations allowed in this state. This list is advisory only and the server state may have changed by the time this field is read by a client.	read only set parameter
current-operations	lists the operations currently in process. This list is advisory only and the server state may have changed by the time this field is read by a client	read only set parameter
patches	Set of host patches	read only set parameter
blobs	Binary data store	read only
memory-free-computed	A conservative estimate of the maximum amount of memory free on a host	read only
ha-statefiles	The UUID(s) of all HA statefiles	read only
ha-network-peers	The UUIDs of all hosts that could host the VMs on this host in case of failure	read only
external-auth-type	Type of external authentication, for example, Active Directory.	read only
external-auth-service-name	The name of the external authentication service	read only
external-auth-configuration	Configuration information for the external authentication service.	read only map parameter

XenServer hosts contain some other objects that also have parameter lists.

CPUs on XenServer hosts have the following parameters:

Parameter Name	Description	Type
uuid	The unique identifier/object reference for the CPU	read only
number	the number of the physical CPU core within the XenServer host	read only
vendor	the vendor string for the CPU name, for example, "GenuineIntel"	read only
speed	The CPU clock speed, in Hz	read only
modelName	the vendor string for the CPU model, for example, "Intel(R) Xeon(TM) CPU 3.00GHz"	read only
stepping	the CPU revision number	read only
flags	the flags of the physical CPU (a decoded version of the features field)	read only
Utilization	the current CPU utilization	read only
host-uuid	the UUID of the host the CPU is in	read only
model	the model number of the physical CPU	read only
family	the physical CPU family number	read only

Crash dumps on XenServer hosts have the following parameters:

Parameter Name	Description	Type
uuid	The unique identifier/object reference for the crashdump	read only
host	XenServer host the crashdump corresponds to	read only
timestamp	Timestamp of the date and time that the crashdump occurred, in the form <i>yyyymmdd-hhmmss-ABC</i> , where <i>ABC</i> is the timezone indicator, for example, GMT	read only
size	size of the crashdump, in bytes	read only

host-backup

`host-backup file-name=<backup_filename> host=<host_name>`

Download a backup of the control domain of the specified XenServer host to the machine that the command is invoked from, and save it there as a file with the name *file-name*.

While the **xe host-backup** command will work if executed on the local host (that is, without a specific hostname specified), do *not* use it this way. Doing so would fill up the control domain partition with the



backup file. The command should *only* be used from a remote off-host machine where you have space to hold the backup file.

host-bugreport-upload

```
host-bugreport-upload [<host-selector>=<host_selector_value>...] [url=<destination_url>]  
[http-proxy=<http_proxy_name>]
```

Generate a fresh bug report (using *xen-bugtool*, with all optional files included) and upload to the Citrix Support ftp site or some other location.

The host(s) on which this operation should be performed are selected using the standard selection mechanism (see [host selectors](#) above). Optional arguments can be any number of the [host selectors](#) listed at the beginning of this section.

Optional parameters are *http-proxy*: use specified http proxy, and *url*: upload to this destination URL. If optional parameters are not used, no proxy server is identified and the destination will be the default Citrix Support ftp site.

host-crashdump-destroy

```
host-crashdump-destroy uuid=<crashdump_uuid>
```

Delete a host crashdump specified by its UUID from the XenServer host.

host-crashdump-upload

```
host-crashdump-upload uuid=<crashdump_uuid>  
[url=<destination_url>]  
[http-proxy=<http_proxy_name>]
```

Upload a crashdump to the Citrix Support ftp site or other location. If optional parameters are not used, no proxy server is identified and the destination will be the default Citrix Support ftp site. Optional parameters are *http-proxy*: use specified http proxy, and *url*: upload to this destination URL.

host-disable

```
host-disable [<host-selector>=<host_selector_value>...]
```

Disables the specified XenServer hosts, which prevents any new VMs from starting on them. This prepares the XenServer hosts to be shut down or rebooted.

The host(s) on which this operation should be performed are selected using the standard selection mechanism (see [host selectors](#) above). Optional arguments can be any number of the [host selectors](#) listed at the beginning of this section.

host-dmesg

```
host-dmesg [<host-selector>=<host_selector_value>...]
```

Get a Xen *dmesg* (the output of the kernel ring buffer) from specified XenServer hosts.

The host(s) on which this operation should be performed are selected using the standard selection mechanism (see [host selectors](#) above). Optional arguments can be any number of the [host selectors](#) listed at the beginning of this section.

host-emergency-management-reconfigure

```
host-emergency-management-reconfigure interface=<uuid_of_management_interface_pif>
```

Reconfigure the management interface of this XenServer host. Use this command only if the XenServer host is in emergency mode, meaning that it is a member in a resource pool whose master has disappeared from the network and could not be contacted for some number of retries.

host-enable

```
host-enable [<host-selector>=<host_selector_value>...]
```

Enables the specified XenServer hosts, which allows new VMs to be started on them.

The host(s) on which this operation should be performed are selected using the standard selection mechanism (see [host selectors](#) above). Optional arguments can be any number of the [host selectors](#) listed at the beginning of this section.

host-evacuate

```
host-evacuate [<host-selector>=<host_selector_value>...]
```

Live migrates all running VMs to other suitable hosts on a pool. The host must first be disabled using the **host-disable** command.

If the evacuated host is the pool master, then another host must be selected to be the pool master. To change the pool master with HA disabled, you need to use the **pool-designate-new-master** command. See [the section called “pool-designate-new-master”](#) for details. With HA enabled, your only option is to shut down the server, which will cause HA to elect a new master at random. See [the section called “host-shutdown”](#).

The host(s) on which this operation should be performed are selected using the standard selection mechanism (see [host selectors](#) above). Optional arguments can be any number of the [host selectors](#) listed at the beginning of this section.

host-forget

```
host-forget uuid=<XenServer_host_UUID>
```

The xapi agent forgets about the specified XenServer host without contacting it explicitly.

Use the **--force** parameter to avoid being prompted to confirm that you really want to perform this operation.

Warning:

Don't use this command if HA is enabled on the pool. Disable HA first, then enable it again after you've forgotten the host.

Tip:

This command is useful if the XenServer host to "forget" is dead; however, if the XenServer host is live and part of the pool, you should use **xe pool-eject** instead.

host-get-system-status

```
host-get-system-status filename=<name_for_status_file>
[entries=<comma_separated_list>] [output=<tar.bz2 | zip>] [<host-selector>=<host_selector_value>...]
```

Download system status information into the specified file. The optional parameter *entries* is a comma-separated list of system status entries, taken from the capabilities XML fragment returned by the **host-get-system-status-capabilities** command. See [the section called “host-get-system-status-capabilities”](#) for



details. If not specified, all system status information is saved in the file. The parameter *output* may be *tar.bz2* (the default) or *zip*; if this parameter is not specified, the file is saved in *tar.bz2* form.

The host(s) on which this operation should be performed are selected using the standard selection mechanism (see [host selectors](#) above).

host-get-system-status-capabilities

host-get-system-status-capabilities [*<host-selector>=<host_selector_value>...*]

Get system status capabilities for the specified host(s). The capabilities are returned as an XML fragment that looks something like this:

```
<?xml version="1.0" ?> <system-status-capabilities>
  <capability content-type="text/plain" default-checked="yes" key="xenserver-logs" \
    max-size="150425200" max-time="-1" min-size="150425200" min-time="-1" \
    pii="maybe"/>
  <capability content-type="text/plain" default-checked="yes" \
    key="xenserver-install" max-size="51200" max-time="-1" min-size="10240" \
    min-time="-1" pii="maybe"/>
  ...
</system-status-capabilities>
```

Each capability entity has a number of attributes.

Attribute	Description
key	A unique identifier for the capability.
content-type	Can be either <i>text/plain</i> or <i>application/data</i> . Indicates whether a UI can render the entries for human consumption.
default-checked	Can be either <i>yes</i> or <i>no</i> . Indicates whether a UI should select this entry by default.
min-size, max-size	Indicates an approximate range for the size, in bytes, of this entry. <i>-1</i> indicates that the size is unimportant.
min-time, max-time	Indicate an approximate range for the time, in seconds, taken to collect this entry. <i>-1</i> indicates the time is unimportant.
pii	<p>Personally identifiable information. Indicates whether the entry would have information that would identify the system owner, or details of their network topology. This is one of:</p> <ul style="list-style-type: none">• <i>no</i>: no PII will be in these entries• <i>yes</i>: PII will likely or certainly be in these entries• <i>maybe</i>: you might wish to audit these entries for PII• <i>if_customized</i> if the files are unmodified, then they will contain no PII, but since we encourage editing of these files, PII may have been introduced by such customization. This is used in particular for the networking scripts in the control domain. <p>Passwords are never to be included in any bug report, regardless of any PII declaration.</p>



The host(s) on which this operation should be performed are selected using the standard selection mechanism (see [host selectors](#) above).

host-is-in-emergency-mode

host-is-in-emergency-mode

Returns *true* if the host the CLI is talking to is currently in emergency mode, *false* otherwise. This CLI command works directly on slave hosts even with no master host present.

host-apply-edition

host-apply-edition [host-uuid=<XenServer_host_UUID>]
[edition=xenserver_edition=<"free"><"advanced"><"enterprise"><"platinum"><"enterprise-xd">]

Assigns a XenServer license to a host server. When you assign a license, XenServer contacts the Citrix License Server and requests the specified type of license. If a license is available, it is then checked out from the license server.

For Citrix XenServer for XenDesktop editions, use <"enterprise-xd">.

For initial licensing configuration, see also **license-server-address** and **license-server-port**.

license-server-address

license-server-address [license-server-address=license_server_address host-
uuid=XenServer<_host_UUID>]

For XenServer Advanced Edition and higher, use to specify either a license server name or IP address.

Run on the pool master before initial use. Specifies the name of the license server the pool is to use. Assigns a XenServer license to a host server. Optionally use with the **license-server-port** command. You only need to set the license server address once and the information is retained until you change it.

license-server-port

license-server-port [license-server-port=license_server_port host-uuid=XenServer<_host_UUID>]

For XenServer Advanced Edition and higher, specifies the port the host is to use to communicate with the Citrix License Server. The default port is 27000, which is the port the license server uses by default for communications with Citrix products. If you changed the port on the Citrix License Server, specify the new port number using this command. Otherwise, you do not need to run this command. For more information about changing port numbers due to conflicts, see the licensing topics in [Citrix eDocs](#).

host-license-add

host-license-add [license-file=<path/license_filename>] [host-uuid=<XenServer_host_UUID>]

For XenServer (free edition), use to parse a local license file and adds it to the specified XenServer host.

Note:

This command only applies to free XenServer. XenServer 5.6 Advanced edition and higher use the **host-apply-edition**, **license-server-address** and **license-server-port** commands. For these editions licensing has changed. They now use the licensing model described in the *Citrix XenServer 5.6 Installation Guide*.



host-license-view

```
host-license-view [host-uuid=<XenServer_host_UUID>]
```

For XenServer (free edition), displays the contents of the XenServer host license.

Note:

This command only applies to free XenServer. XenServer 5.6 Advanced edition and higher use the **host-apply-edition**, **license-server-address** and **license-server-port** commands. For these editions licensing has changed. They now use the licensing model described in the *Citrix XenServer 5.6 Installation Guide*.

host-logs-download

```
host-logs-download [file-name=<logfile_name>] [<host-selector>=<host_selector_value>...]
```

Download a copy of the logs of the specified XenServer hosts. The copy is saved by default in a time-stamped file named `hostname-yyyy-mm-dd T hh:mm:ssZ.tar.gz`. You can specify a different filename using the optional parameter *file-name*.

The host(s) on which this operation should be performed are selected using the standard selection mechanism (see [host selectors](#) above). Optional arguments can be any number of the [host selectors](#) listed at the beginning of this section.

While the **xe host-logs-download** command will work if executed on the local host (that is, without a specific hostname specified), do *not* use it this way. Doing so will clutter the control domain partition with the copy of the logs. The command should *only* be used from a remote off-host machine where you have space to hold the copy of the logs.

host-management-disable

```
host-management-disable
```

Disables the host agent listening on an external management network interface and disconnects all connected API clients (such as the XenCenter). Operates directly on the XenServer host the CLI is connected to, and is not forwarded to the pool master if applied to a member XenServer host.

Warning:

Be extremely careful when using this CLI command off-host, since once it is run it will not be possible to connect to the control domain remotely over the network to re-enable it.

host-management-reconfigure

```
host-management-reconfigure [interface=<device> ] | [pif-uuid=<uuid> ]
```

Reconfigures the XenServer host to use the specified network interface as its management interface, which is the interface that is used to connect to the XenCenter. The command rewrites the `MANAGEMENT_INTERFACE` key in `/etc/xensource-inventory`.

If the device name of an interface (which must have an IP address) is specified, the XenServer host will immediately rebind. This works both in normal and emergency mode.

If the UUID of a PIF object is specified, the XenServer host determines which IP address to rebind to itself. It must not be in emergency mode when this command is executed.

Warning:



Be careful when using this CLI command off-host and ensure that you have network connectivity on the new interface. Use **xe pif-reconfigure** to set one up first. Otherwise, subsequent CLI commands will reach the XenServer host.

host-power-on

```
host-power-on [host=<host_uuid> ]
```

Turns on power on XenServer hosts with Host Power On functionality enabled. Before using this command, **host-set-power-on** must be enabled on the host.

host-set-power-on

```
host-set-power-on {host=<host uuid> {power-on-mode=<""> <"wake-on-lan"> <"iLO"> <"DRAC">  
<"custom"> } | [power-on-config=<"power_on_ip"><"power_on_user"><"power_on_password_secret">]}
```

Use to enable Host Power On functionality on XenServer hosts that are compatible with remote power solutions. Workload Balancing requires Host Power On functionality is enabled for it to turn off underused hosts in Maximum Density mode. When using the **host-set-power-on** command, you must specify the type of power management solution on the host (that is, the **<power-on-mode>**). Then specify configuration options using the **<power-on-config>** argument and its associated key-value pairs. To use the secrets feature to store your password, specify the key "power_on_password_secret".

host-reboot

```
host-reboot [<host-selector>=<host_selector_value>...]
```

Reboot the specified XenServer hosts. The specified XenServer hosts must be disabled first using the **xe host-disable** command, otherwise a **HOST_IN_USE** error message is displayed.

The host(s) on which this operation should be performed are selected using the standard selection mechanism (see [host selectors](#) above). Optional arguments can be any number of the [host selectors](#) listed at the beginning of this section.

If the specified XenServer hosts are members of a pool, the loss of connectivity on shutdown will be handled and the pool will recover when the XenServer hosts returns. If you shut down a pool member, other members and the master will continue to function. If you shut down the master, the pool will be out of action until the master is rebooted and back on line (at which point the members will reconnect and synchronize with the master) or until you make one of the members into the master.

host-restore

```
host-restore [file-name=<backup_filename>] [<host-selector>=<host_selector_value>...]
```

Restore a backup named *file-name* of the XenServer host control software. Note that the use of the word "restore" here does not mean a full restore in the usual sense, it merely means that the compressed backup file has been uncompressed and unpacked onto the secondary partition. After you've done a **xe host-restore**, you have to boot the Install CD and use its **Restore from Backup** option.

The host(s) on which this operation should be performed are selected using the standard selection mechanism (see [host selectors](#) above). Optional arguments can be any number of the [host selectors](#) listed at the beginning of this section.

host-set-hostname-live

```
host-set-hostname host-uuid=<uuid_of_host> hostname=<new_hostname>
```



Change the hostname of the XenServer host specified by *host-uuid*. This command persistently sets both the hostname in the control domain database and the actual Linux hostname of the XenServer host. Note that *hostname* is *not* the same as the value of the *name_label* field.

host-shutdown

```
host-shutdown [<host-selector>=<host_selector_value>...]
```

Shut down the specified XenServer hosts. The specified XenServer hosts must be disabled first using the **xe host-disable** command, otherwise a `HOST_IN_USE` error message is displayed.

The host(s) on which this operation should be performed are selected using the standard selection mechanism (see [host selectors](#) above). Optional arguments can be any number of the [host selectors](#) listed at the beginning of this section.

If the specified XenServer hosts are members of a pool, the loss of connectivity on shutdown will be handled and the pool will recover when the XenServer hosts returns. If you shut down a pool member, other members and the master will continue to function. If you shut down the master, the pool will be out of action until the master is rebooted and back on line, at which point the members will reconnect and synchronize with the master, or until one of the members is made into the master. If HA is enabled for the pool, one of the members will be made into a master automatically. If HA is disabled, you must manually designate the desired server as master with the **pool-designate-new-master** command. See [the section called “pool-designate-new-master”](#).

host-syslog-reconfigure

```
host-syslog-reconfigure [<host-selector>=<host_selector_value>...]
```

Reconfigure the `syslog` daemon on the specified XenServer hosts. This command applies the configuration information defined in the host *logging* parameter.

The host(s) on which this operation should be performed are selected using the standard selection mechanism (see [host selectors](#) above). Optional arguments can be any number of the [host selectors](#) listed at the beginning of this section.

Log commands

Commands for working with logs.

log-get-keys

```
log-get-keys
```

List the keys of all of the logging subsystems.

log-reopen

```
log-reopen
```

Reopen all loggers. Use this command for rotating log files.

log-set-output

```
log-set-output output=nil | stderr | file:<filename> | syslog:<sysloglocation> [key=<key>] [level= debug  
| info | warning | error]
```



Set the output of the specified logger. Log messages are filtered by the subsystem in which they originated and the log level of the message. For example, send debug logging messages from the storage manager to a file by running the following command:

```
xe log-set-output key=sm level=debug output=<file:/tmp/sm.log>
```

The optional parameter *key* specifies the particular logging subsystem. If this parameter is not set, it will default to all logging subsystems.

The optional parameter *level* specifies the logging level. Valid values are:

- debug
- info
- warning
- error

Message commands

Commands for working with messages. Messages are created to notify users of significant events, and are displayed in XenCenter as system alerts.

Message parameters

Parameter Name	Description	Type
uuid	The unique identifier/object reference for the message	read only
name	The unique name of the message	read only
priority	The message priority. Higher numbers indicate greater priority	read only
class	The message class, for example VM.	read only
obj-uuid	The uuid of the affected object.	read only
timestamp	The time that the message was generated.	read only
body	The message content.	read only

message-create

```
message-create name=<message_name> body=<message_text> [[host-uuid=<uuid_of_host>] | [sr-uuid=<uuid_of_sr>] | [vm-uuid=<uuid_of_vm>] | [pool-uuid=<uuid_of_pool>]]
```

Creates a new message.

message-list

```
message-list
```

Lists all messages, or messages that match the specified standard selectable parameters.

Network commands

Commands for working with networks.

The network objects can be listed with the standard object listing command (**xe network-list**), and the parameters manipulated with the standard parameter commands. See [the section called “Low-level param commands”](#) for details.

Network parameters

Networks have the following parameters:

Parameter Name	Description	Type
uuid	The unique identifier/object reference for the network	read only
name-label	The name of the network	read write
name-description	The description text of the network	read write
VIF-uuids	A list of unique identifiers of the VIFs (virtual network interfaces) that are attached from VMs to this network	read only set parameter
PIF-uuids	A list of unique identifiers of the PIFs (physical network interfaces) that are attached from XenServer hosts to this network	read only set parameter
bridge	name of the bridge corresponding to this network on the local XenServer host	read only
other-config:static-routes	comma-separated list of <subnet>/<netmask>/<gateway> formatted entries specifying the gateway address through which to route subnets. For example, setting other-config:static-routes to 172.16.0.0/15/192.168.0.3,172.18.0.0/16/192.168.0.4 causes traffic on 172.16.0.0/15 to be routed over 192.168.0.3 and traffic on 172.18.0.0/16 to be routed over 192.168.0.4.	read write
other-config:ethtool-autoneg	set to no to disable autonegotiation of the physical interface or bridge. Default is yes .	read write
other-config:ethtool-rx	set to on to enable receive checksum, off to disable	read write
other-config:ethtool-tx	set to on to enable transmit checksum, off to disable	read write

Parameter Name	Description	Type
other-config:ethtool-sg	set to <code>on</code> to enable scatter gather, <code>off</code> to disable	read write
other-config:ethtool-tso	set to <code>on</code> to enable tcp segmentation offload, <code>off</code> to disable	read write
other-config:ethtool-ufo	set to <code>on</code> to enable UDP fragment offload, <code>off</code> to disable	read write
other-config:ethtool-gso	set to <code>on</code> to enable generic segmentation offload, <code>off</code> to disable	read write
blobs	Binary data store	read only

network-create

`network-create name-label=<name_for_network> [name-description=<descriptive_text>]`

Creates a new network.

network-destroy

`network-destroy uuid=<network_uuid>`

Destroys an existing network.

Patch (update) commands

Commands for working with XenServer host patches (updates). These are for the standard non-OEM editions of XenServer for commands relating to updating the OEM edition of XenServer, see [the section called “Update commands”](#) for details.

The patch objects can be listed with the standard object listing command (**`xe patch-list`**), and the parameters manipulated with the standard parameter commands. See [the section called “Low-level param commands”](#) for details.

Patch parameters

Patches have the following parameters:

Parameter Name	Description	Type
uuid	The unique identifier/object reference for the patch	read only
host-uuid	The unique identifier for the XenServer host to query	read only
name-label	The name of the patch	read only
name-description	The description string of the patch	read only

Parameter Name	Description	Type
applied	Whether or not the patch has been applied; true or false	read only
size	Whether or not the patch has been applied; true or false	read only

patch-apply

patch-apply uuid=<patch_file_uuid>

Apply the specified patch file.

patch-clean

patch-clean uuid=<patch_file_uuid>

Delete the specified patch file from the XenServer host.

patch-pool-apply

patch-pool-apply uuid=<patch_uuid>

Apply the specified patch to all XenServer hosts in the pool.

patch-precheck

patch-precheck uuid=<patch_uuid> host-uuid=<host_uuid>

Run the prechecks contained within the specified patch on the specified XenServer host.

patch-upload

patch-upload file-name=<patch_filename>

Upload a specified patch file to the XenServer host. This prepares a patch to be applied. On success, the UUID of the uploaded patch is printed out. If the patch has previously been uploaded, a PATCH_ALREADY_EXISTS error is returned instead and the patch is not uploaded again.

PBD commands

Commands for working with PBDs (Physical Block Devices). These are the software objects through which the XenServer host accesses storage repositories (SRs).

The PBD objects can be listed with the standard object listing command (**xe pbd-list**), and the parameters manipulated with the standard parameter commands. See [the section called “Low-level param commands”](#) for details.

PBD parameters

PBDs have the following parameters:

Parameter Name	Description	Type
uuid	The unique identifier/object reference for the PBD.	read only

Parameter Name	Description	Type
sr-uuid	the storage repository that the PBD points to	read only
device-config	additional configuration information that is provided to the SR-backend-driver of a host	read only map parameter
currently-attached	True if the SR is currently attached on this host, False otherwise	read only
host-uuid	UUID of the physical machine on which the PBD is available	read only
host	The host field is deprecated. Use host_uuid instead.	read only
other-config	Additional configuration information.	read/write map parameter

pbd-create

```
pbd-create host-uuid=<uuid_of_host>
sr-uuid=<uuid_of_sr>
[device-config:key=<corresponding_value>...]
```

Create a new PBD on a XenServer host. The read-only *device-config* parameter can only be set on creation.

To add a mapping of 'path' -> '/tmp', the command line should contain the argument *device-config:path=/tmp*

For a full list of supported device-config key/value pairs on each SR type see [Storage](#).

pbd-destroy

```
pbd-destroy uuid=<uuid_of_pbd>
```

Destroy the specified PBD.

pbd-plug

```
pbd-plug uuid=<uuid_of_pbd>
```

Attempts to plug in the PBD to the XenServer host. If this succeeds, the referenced SR (and the VDIs contained within) should then become visible to the XenServer host.

pbd-unplug

```
pbd-unplug uuid=<uuid_of_pbd>
```

Attempt to unplug the PBD from the XenServer host.

PIF commands

Commands for working with PIFs (objects representing the physical network interfaces).



The PIF objects can be listed with the standard object listing command (**xe pif-list**), and the parameters manipulated with the standard parameter commands. See [the section called “Low-level param commands”](#) for details.

PIF parameters

PIFs have the following parameters:

Parameter Name	Description	Type
uuid	the unique identifier/object reference for the PIF	read only
device	machine-readable name of the interface (for example, eth0)	read only
MAC	the MAC address of the PIF	read only
other-config	Additional PIF configuration name:value pairs.	read/write map parameter
physical	if true, the PIF points to an actual physical network interface	read only
currently-attached	is the PIF currently attached on this host? true or false	read only
MTU	Maximum Transmission Unit of the PIF in bytes.	read only
VLAN	VLAN tag for all traffic passing through this interface; -1 indicates no VLAN tag is assigned	read only
bond-master-of	the UUID of the bond this PIF is the master of (if any)	read only
bond-slave-of	the UUID of the bond this PIF is the slave of (if any)	read only
management	is this PIF designated to be a management interface for the control domain	read only
network-uuid	the unique identifier/object reference of the virtual network to which this PIF is connected	read only
network-name-label	the name of the of the virtual network to which this PIF is connected	read only
host-uuid	the unique identifier/object reference of the XenServer host to which this PIF is connected	read only

Parameter Name	Description	Type
host-name-label	the name of the XenServer host to which this PIF is connected	read only
IP-configuration-mode	type of network address configuration used; DHCP or static	read only
IP	IP address of the PIF, defined here if IP-configuration-mode is static; undefined if DHCP	read only
netmask	Netmask of the PIF, defined here if IP-configuration-mode is static; undefined if supplied by DHCP	read only
gateway	Gateway address of the PIF, defined here if IP-configuration-mode is static; undefined if supplied by DHCP	read only
DNS	DNS address of the PIF, defined here if IP-configuration-mode is static; undefined if supplied by DHCP	read only
io_read_kbs	average read rate in kB/s for the device	read only
io_write_kbs	average write rate in kB/s for the device	read only
carrier	link state for this device	read only
vendor-id	the ID assigned to NIC's vendor	read only
vendor-name	the NIC vendor's name	read only
device-id	the ID assigned by the vendor to this NIC model	read only
device-name	the name assigned by the vendor to this NIC model	read only
speed	data transfer rate of the NIC	read only
duplex	duplexing mode of the NIC; full or half	read only
pci-bus-path	PCI bus path address	read only
other-config:ethtool-speed	sets the speed of connection in Mbps	read write
other-config:ethtool-autoneg	set to <code>no</code> to disable autonegotiation of the physical interface or bridge. Default is <code>yes</code> .	read write
other-config:ethtool-duplex	Sets duplexing capability of the PIF, either <code>full</code> or <code>half</code> .	read write

Parameter Name	Description	Type
other-config:ethtool-rx	set to <code>on</code> to enable receive checksum, <code>off</code> to disable	read write
other-config:ethtool-tx	set to <code>on</code> to enable transmit checksum, <code>off</code> to disable	read write
other-config:ethtool-sg	set to <code>on</code> to enable scatter gather, <code>off</code> to disable	read write
other-config:ethtool-tso	set to <code>on</code> to enable tcp segmentation offload, <code>off</code> to disable	read write
other-config:ethtool-ufo	set to <code>on</code> to enable udp fragment offload, <code>off</code> to disable	read write
other-config:ethtool-gso	set to <code>on</code> to enable generic segmentation offload, <code>off</code> to disable	read write
other-config:domain	comma-separated list used to set the DNS search path	read write
other-config:bond-miimon	interval between link liveness checks, in milliseconds	read write
other-config:bond-downdelay	number of milliseconds to wait after link is lost before really considering the link to have gone. This allows for transient link loss	read write
other-config:bond-updelay	number of milliseconds to wait after the link comes up before really considering it up. Allows for links flapping up. Default is 31s to allow for time for switches to begin forwarding traffic.	read write
disallow-unplug	True if this PIF is a dedicated storage NIC, false otherwise	read/write

Note:

Changes made to the `other-config` fields of a PIF will only take effect after a reboot. Alternately, use the **xe pif-unplug** and **xe pif-plug** commands to cause the PIF configuration to be rewritten.

pif-forget

`pif-forget uuid=<uuid_of_pif>`

Destroy the specified PIF object on a particular host.

pif-introduce

`pif-introduce host-uuid=<UUID of XenServer host> mac=<mac_address_for_pif> device=<machine-readable name of the interface (for example, eth0)>`

Create a new PIF object representing a physical interface on the specified XenServer host.



pif-plug

```
pif-plug uuid=<uuid_of_pif>
```

Attempt to bring up the specified physical interface.

pif-reconfigure-ip

```
pif-reconfigure-ip uuid=<uuid_of_pif> [ mode=<dhcp> | mode=<static> ]  
gateway=<network_gateway_address> IP=<static_ip_for_this_pif>  
netmask=<netmask_for_this_pif> [DNS=<dns_address>]
```

Modify the IP address of the PIF. For static IP configuration, set the *mode* parameter to *static*, with the *gateway*, *IP*, and *netmask* parameters set to the appropriate values. To use DHCP, set the *mode* parameter to *DHCP* and leave the static parameters undefined.

Note:

Using static IP addresses on physical network interfaces connected to a port on a switch using Spanning Tree Protocol with STP Fast Link turned off (or unsupported) results in a period during which there is no traffic.

pif-scan

```
pif-scan host-uuid=<UUID of XenServer host>
```

Scan for new physical interfaces on a XenServer host.

pif-unplug

```
pif-unplug uuid=<uuid_of_pif>
```

Attempt to bring down the specified physical interface.

Pool commands

Commands for working with pools. A *pool* is an aggregate of one or more XenServer hosts. A pool uses one or more shared storage repositories so that the VMs running on one XenServer host in the pool can be migrated in near-real time (while still running, without needing to be shut down and brought back up) to another XenServer host in the pool. Each XenServer host is really a pool consisting of a single member by default. When a XenServer host is joined to a pool, it is designated as a member, and the pool it has joined becomes the master for the pool.

The singleton pool object can be listed with the standard object listing command (**xe pool-list**), and its parameters manipulated with the standard parameter commands. See [the section called “Low-level param commands”](#) for details.

Pool parameters

Pools have the following parameters:

Parameter Name	Description	Type
uuid	the unique identifier/object reference for the pool	read only
name-label	the name of the pool	read/write
name-description	the description string of the pool	read/write

Parameter Name	Description	Type
master	the unique identifier/object reference of XenServer host designated as the pool's master	read only
default-SR	the unique identifier/object reference of the default SR for the pool	read/write
crash-dump-SR	the unique identifier/object reference of the SR where any crash dumps for pool members are saved	read/write
suspend-image-SR	the unique identifier/object reference of the SR where suspended VMs on pool members are saved	read/write
other-config	a list of key/value pairs that specify additional configuration parameters for the pool	read/write map parameter
supported-sr-types	SR types that can be used by this pool	read only
ha-enabled	True if HA is enabled for the pool, false otherwise	read only
ha-configuration	reserved for future use.	read only
ha-statefiles	lists the UUIDs of the VDIs being used by HA to determine storage health	read only
ha-host-failures-to-tolerate	the number of host failures to tolerate before sending a system alert	read/write
ha-plan-exists-for	the number of hosts failures that can actually be handled, according to the calculations of the HA algorithm	read only
ha-allow-overcommit	True if the pool is allowed to be overcommitted, False otherwise	read/write
ha-overcommitted	True if the pool is currently overcommitted	read only

Parameter Name	Description	Type
blobs	binary data store	read only
wlb-url	Path to the WLB server	read only
wlb-username	Name of the user of the WLB service	read only
wlb-enabled	True is WLB is enabled	read/write
wlb-verify-cert	True if there is a certificate to verify	read/write

pool-designate-new-master

`pool-designate-new-master host-uuid=<UUID of member XenServer host to become new master>`

Instruct the specified member XenServer host to become the master of an existing pool. This performs an orderly hand over of the role of master host to another host in the resource pool. This command only works when the current master is online, and is not a replacement for the emergency mode commands listed below.

pool-dump-database

`pool-dump-database file-name=<filename_to_dump_database_into_(on_client)>`

Download a copy of the entire pool database and dump it into a file on the client.

pool-eject

`pool-eject host-uuid=<UUID of XenServer host to eject>`

Instruct the specified XenServer host to leave an existing pool.

pool-emergency-reset-master

`pool-emergency-reset-master master-address=<address of the pool's master XenServer host>`

Instruct a slave member XenServer host to reset its master address to the new value and attempt to connect to it. This command should not be run on master hosts.

pool-emergency-transition-to-master

`pool-emergency-transition-to-master`

Instruct a member XenServer host to become the pool master. This command is only accepted by the XenServer host if it has transitioned to emergency mode, meaning it is a member of a pool whose master has disappeared from the network and could not be contacted for some number of retries.

Note that this command may cause the password of the host to reset if it has been modified since joining the pool (see [the section called "User commands"](#)).

pool-ha-enable

`pool-ha-enable heartbeat-sr-uuids=<SR_UUID_of_the_Heartbeat_SR>`

Enable High Availability on the resource pool, using the specified SR UUID as the central storage heartbeat repository.



pool-ha-disable

pool-ha-disable

Disables the High Availability functionality on the resource pool.

pool-join

pool-join master-address=<address> master-username=<username> master-password=<password>

Instruct a XenServer host to join an existing pool.

pool-recover-slaves

pool-recover-slaves

Instruct the pool master to try and reset the master address of all members currently running in emergency mode. This is typically used after **pool-emergency-transition-to-master** has been used to set one of the members as the new master.

pool-restore-database

pool-restore-database file-name=<filename_to_restore_from_(on_client)> [dry-run=<true | false>]

Upload a database backup (created with **pool-dump-database**) to a pool. On receiving the upload, the master will restart itself with the new database.

There is also a *dry run* option, which allows you to check that the pool database can be restored without actually perform the operation. By default, `dry-run` is set to false.

pool-sync-database

pool-sync-database

Force the pool database to be synchronized across all hosts in the resource pool. This is not necessary in normal operation since the database is regularly automatically replicated, but can be useful for ensuring changes are rapidly replicated after performing a significant set of CLI operations.

Storage Manager commands

Commands for controlling Storage Manager plugins.

The storage manager objects can be listed with the standard object listing command (**xe sm-list**), and the parameters manipulated with the standard parameter commands. See [the section called “Low-level param commands”](#) for details.

SM parameters

SMs have the following parameters:

Parameter Name	Description	Type
uuid	the unique identifier/object reference for the SM plugin	read only
name-label	the name of the SM plugin	read only
name-description	the description string of the SM plugin	read only
type	the SR type that this plugin connects to	read only

Parameter Name	Description	Type
vendor	name of the vendor who created this plugin	read only
copyright	copyright statement for this SM plugin	read only
required-api-version	minimum SM API version required on the XenServer host	read only
configuration	names and descriptions of device configuration keys	read only
capabilities	capabilities of the SM plugin	read only
driver-filename	the filename of the SR driver.	read only

SR commands

Commands for controlling SRs (storage repositories).

The SR objects can be listed with the standard object listing command (**xe sr-list**), and the parameters manipulated with the standard parameter commands. See [the section called “Low-level param commands”](#) for details.

SR parameters

SRs have the following parameters:

Parameter Name	Description	Type
uuid	the unique identifier/object reference for the SR	read only
name-label	the name of the SR	read/write
name-description	the description string of the SR	read/write
allowed-operations	list of the operations allowed on the SR in this state	read only set parameter
current-operations	list of the operations that are currently in progress on this SR	read only set parameter
VDIs	unique identifier/object reference for the virtual disks in this SR	read only set parameter
PBDs	unique identifier/object reference for the PBDs attached to this SR	read only set parameter
physical-Utilization	physical space currently utilized on this SR, in bytes. Note that for sparse disk formats, physical utilization may be less than virtual allocation	read only

Parameter Name	Description	Type
physical-size	total physical size of the SR, in bytes	read only
type	type of the SR, used to specify the SR backend driver to use	read only
content-type	the type of the SR's content. Used to distinguish ISO libraries from other SRs. For storage repositories that store a library of ISOs, the <i>content-type</i> must be set to <i>iso</i> . In other cases, Citrix recommends that this be set either to empty, or the string <i>user</i> .	read only
shared	True if this SR is capable of being shared between multiple XenServer hosts; False otherwise	read/write
other-config	list of key/value pairs that specify additional configuration parameters for the SR .	read/write map parameter
host	The storage repository host name	read only
virtual-allocation	sum of virtual-size values of all VDIs in this storage repository (in bytes)	read only
sm-config	SM dependent data	read only map parameter
blobs	binary data store	read only

sr-create

```
sr-create name-label=<name> physical-size=<size> type=<type>
content-type=<content_type> device-config:<config_name>=<value>
[host-uuid=<XenServer host UUID>] [shared=<true | false>]
```

Creates an SR on the disk, introduces it into the database, and creates a PBD attaching the SR to a XenServer host. If *shared* is set to *true*, a PBD is created for each XenServer host in the pool; if *shared* is not specified or set to *false*, a PBD is created only for the XenServer host specified with *host-uuid*.

The exact *device-config* parameters differ depending on the device *type*. See [Storage](#) for details of these parameters across the different storage backends.

sr-destroy

```
sr-destroy uuid=<sr_uuid>
```

Destroys the specified SR on the XenServer host.

sr-forget

```
sr-forget uuid=<sr_uuid>
```

The *xapi* agent forgets about a specified SR on the XenServer host, meaning that the SR is detached and you cannot access VDIs on it, but it remains intact on the source media (the data is not lost).



sr-introduce

```
sr-introduce name-label=<name>  
physical-size=<physical_size>  
type=<type>  
content-type=<content_type>  
uuid=<sr_uuid>
```

Just places an SR record into the database. The *device-config* parameters are specified by *device-config*:<parameter_key>=<parameter_value>, for example:

```
xe sr-introduce device-config:<device>=</dev/sdb1>
```

Note:

This command is never used in normal operation. It is an advanced operation which might be useful if an SR needs to be reconfigured as shared after it was created, or to help recover from various failure scenarios.

sr-probe

```
sr-probe type=<type> [host-uuid=<uuid_of_host>] [device-config:<config_name>=<value>]
```

Performs a backend-specific scan, using the provided *device-config* keys. If the *device-config* is complete for the SR backend, then this will return a list of the SRs present on the device, if any. If the *device-config* parameters are only partial, then a backend-specific scan will be performed, returning results that will guide you in improving the remaining *device-config* parameters. The scan results are returned as backend-specific XML, printed out on the CLI.

The exact *device-config* parameters differ depending on the device *type*. See [Storage](#) for details of these parameters across the different storage backends.

sr-scan

```
sr-scan uuid=<sr_uuid>
```

Force an SR scan, syncing the *xapi* database with VDIs present in the underlying storage substrate.

Task commands

Commands for working with long-running asynchronous tasks. These are tasks such as starting, stopping, and suspending a Virtual Machine, which are typically made up of a set of other atomic subtasks that together accomplish the requested operation.

The task objects can be listed with the standard object listing command (**xe task-list**), and the parameters manipulated with the standard parameter commands. See [the section called “Low-level param commands”](#) for details.

Task parameters

Tasks have the following parameters:

Parameter Name	Description	Type
uuid	the unique identifier/object reference for the Task	read only
name-label	the name of the Task	read only

Parameter Name	Description	Type
name-description	the description string of the Task	read only
resident-on	the unique identifier/object reference of the host on which the task is running	read only
status	current status of the Task	read only
progress	if the Task is still pending, this field contains the estimated percentage complete, from 0. to 1. If the Task has completed, successfully or unsuccessfully, this should be 1.	read only
type	if the Task has successfully completed, this parameter contains the type of the encoded result - that is, the name of the class whose reference is in the result field; otherwise, this parameter's value is undefined	read only
result	if the Task has completed successfully, this field contains the result value, either Void or an object reference; otherwise, this parameter's value is undefined	read only
error_info	if the Task has failed, this parameter contains the set of associated error strings; otherwise, this parameter's value is undefined	read only
allowed_operations	list of the operations allowed in this state	read only
created	time the task has been created	read only
finished	time task finished (i.e. succeeded or failed). If task-status is pending, then the value of this field has no meaning	read only
subtask_of	contains the UUID of the tasks this task is a sub-task of	read only
subtasks	contains the UUID(s) of all the subtasks of this task	read only

task-cancel

task-cancel [uuid=<task_uuid>]

Direct the specified Task to cancel and return.

Template commands

Commands for working with VM templates.

Templates are essentially VMs with the *is-a-template* parameter set to *true*. A template is a "gold image" that contains all the various configuration settings to instantiate a specific VM. XenServer ships with a base set of templates, which range from generic "raw" VMs that can boot an OS vendor installation CD (RHEL, CentOS, SLES, Windows) to complete pre-configured OS instances (the "Demo Linux VM" template). With XenServer you can create VMs, configure them in standard forms for your particular needs, and save a copy of them as templates for future use in VM deployment.

The template objects can be listed with the standard object listing command (**xe template-list**), and the parameters manipulated with the standard parameter commands. See [the section called "Low-level param commands"](#) for details.

Template parameters

Templates have the following parameters:

Parameter Name	Description	Type
uuid	the unique identifier/object reference for the template	read only
name-label	the name of the template	read/write
name-description	the description string of the template	read/write
user-version	string for creators of VMs and templates to put version information	read/write
is-a-template	<p>true if this is a template. Template VMs can never be started, they are used only for cloning other VMs</p> <p>Note that setting <i>is-a-template</i> using the CLI is not supported.</p>	read/write
is-control-domain	true if this is a control domain (domain 0 or a driver domain)	read only
power-state	current power state; always <i>halted</i> for a template	read only
power-state	current power state; always <i>halted</i> for a template	read only
memory-dynamic-max	<p>dynamic maximum memory in bytes. Currently unused, but if changed the following constraint must be obeyed:</p> <pre>memory_static_max >= memory_dynamic_max >= memory_dynamic_min >= memory_static_min.</pre>	read/write

Parameter Name	Description	Type
memory-dynamic-min	dynamic minimum memory in bytes. Currently unused, but if changed the same constraints for <i>memory-dynamic-max</i> must be obeyed.	read/write
memory-static-max	statically-set (absolute) maximum memory in bytes. This is the main value used to determine the amount of memory assigned to a VM.	read/write
memory-static-min	statically-set (absolute) minimum memory in bytes. This represents the absolute minimum memory, and <i>memory-static-min</i> must be less than <i>memory-static-max</i> . This value is currently unused in normal operation, but the previous constraint must be obeyed.	read/write
suspend-VDI-uuid	the VDI that a suspend image is stored on (has no meaning for a template)	read only

Parameter Name	Description	Type
VCPUs-params	<p>configuration parameters for the selected VCPU policy.</p> <p>You can tune a VCPU's pinning with</p> <pre>xe vm-param-set \ uuid=<vm_uuid> \ VCPUs-params:mask=1,2,3</pre> <p>A VM created from this template will then run on physical CPUs 1, 2, and 3 only.</p> <p>You can also tune the VCPU priority (xen scheduling) with the <i>cap</i> and <i>weight</i> parameters; for example</p> <pre>xe vm-param-set \ uuid=<vm_uuid> \ VCPUs-params:weight=512 xe vm-param-set \ uuid=<vm_uuid> \ VCPUs-params:cap=100</pre> <p>A VM based on this template with a weight of 512 will get twice as much CPU as a domain with a weight of 256 on a contended XenServer host. Legal weights range from 1 to 65535 and the default is 256.</p> <p>The cap optionally fixes the maximum amount of CPU a VM based on this template will be able to consume, even if the XenServer host has idle CPU cycles. The cap is expressed in percentage of one physical CPU: 100 is 1 physical CPU, 50 is half a CPU, 400 is 4 CPUs, etc. The default, 0, means there is no upper cap.</p>	read/write map parameter
VCPUs-max	maximum number of VCPUs	read/write
VCPUs-at-startup	boot number of VCPUs	read/write

Parameter Name	Description	Type
actions-after-crash	action to take if a VM based on this template crashes	read/write
console-uuids	virtual console devices	read only set parameter
platform	platform-specific configuration	read/write map parameter
allowed-operations	list of the operations allowed in this state	read only set parameter
current-operations	list of the operations that are currently in progress on this template	read only set parameter
allowed-VBD-devices	list of VBD identifiers available for use, represented by integers of the range 0-15. This list is informational only, and other devices may be used (but may not work).	read only set parameter
allowed-VIF-devices	list of VIF identifiers available for use, represented by integers of the range 0-15. This list is informational only, and other devices may be used (but may not work).	read only set parameter
HVM-boot-policy	the boot policy for HVM guests. Either <code>BIOS Order</code> or an empty string.	read/write
HVM-boot-params	the <code>order</code> key controls the HVM guest boot order, represented as a string where each character is a boot method: <code>d</code> for the CD/DVD, <code>c</code> for the root disk, and <code>n</code> for network PXE boot. The default is <code>dc</code> .	read/write map parameter
PV-kernel	path to the kernel	read/write
PV-ramdisk	path to the initrd	read/write
PV-args	string of kernel command line arguments	read/write
PV-legacy-args	string of arguments to make legacy VMs based on this template boot	read/write

Parameter Name	Description	Type
PV-bootloader	name of or path to bootloader	read/write
PV-bootloader-args	string of miscellaneous arguments for the bootloader	read/write
last-boot-CPU-flags	describes the CPU flags on which a VM based on this template was last booted; not populated for a template	read only
resident-on	the XenServer host on which a VM based on this template is currently resident; appears as <code><not in database></code> for a template	read only
affinity	a XenServer host which a VM based on this template has preference for running on; used by the xe vm-start command to decide where to run the VM	read/write
other-config	list of key/value pairs that specify additional configuration parameters for the template	read/write map parameter
start-time	timestamp of the date and time that the metrics for a VM based on this template were read, in the form <code>yyyymmddThh:mm:ss z</code> , where <code>z</code> is the single-letter military timezone indicator, for example, <code>Z</code> for UTC (GMT); set to 1 Jan 1970 Z (beginning of Unix/POSIX epoch) for a template	read only
install-time	timestamp of the date and time that the metrics for a VM based on this template were read, in the form <code>yyyymmddThh:mm:ss z</code> , where <code>z</code> is the single-letter military timezone indicator, for example, <code>Z</code> for UTC (GMT); set to 1 Jan 1970 Z (beginning of Unix/POSIX epoch) for a template	read only

Parameter Name	Description	Type
memory-actual	the actual memory being used by a VM based on this template; 0 for a template	read only
VCPUs-number	the number of virtual CPUs assigned to a VM based on this template; 0 for a template	read only
VCPUs-Utilization	list of virtual CPUs and their weight	read only map parameter
os-version	the version of the operating system for a VM based on this template; appears as <not in database> for a template	read only map parameter
PV-drivers-version	the versions of the paravirtualized drivers for a VM based on this template; appears as <not in database> for a template	read only map parameter
PV-drivers-up-to-date	flag for latest version of the paravirtualized drivers for a VM based on this template; appears as <not in database> for a template	read only
memory	memory metrics reported by the agent on a VM based on this template; appears as <not in database> for a template	read only map parameter
disks	disk metrics reported by the agent on a VM based on this template; appears as <not in database> for a template	read only map parameter
networks	network metrics reported by the agent on a VM based on this template; appears as <not in database> for a template	read only map parameter
other	other metrics reported by the agent on a VM based on this template; appears as <not in database> for a template	read only map parameter

Parameter Name	Description	Type
guest-metrics-last-updated	timestamp when the last write to these fields was performed by the in-guest agent, in the form <code>yyyymmddThh:mm:ssz</code> , where <code>z</code> is the single-letter military timezone indicator, for example, <code>Z</code> for UTC (GMT)	read only
actions-after-shutdown	action to take after the VM has shutdown	read/write
actions-after-reboot	action to take after the VM has rebooted	read/write
possible-hosts	list of hosts that could potentially host the VM	read only
HVM-shadow-multiplier	multiplier applied to the amount of shadow that will be made available to the guest	read/write
dom-id	domain ID (if available, <code>-1</code> otherwise)	read only
recommendations	XML specification of recommended values and ranges for properties of this VM	read only
xenstore-data	data to be inserted into the xenstore tree (<code>/local/domain/<domid>/vm-data</code>) after the VM is created.	read/write map parameter
is-a-snapshot	True if this template is a VM snapshot	read only
snapshot_of	the UUID of the VM that this template is a snapshot of	read only
snapshots	the UUID(s) of any snapshots that have been taken of this template	read only
snapshot_time	the timestamp of the most recent VM snapshot taken	read only
memory-target	the target amount of memory set for this template	read only

Parameter Name	Description	Type
blocked-operations	lists the operations that cannot be performed on this template	read/write map parameter
last-boot-record	record of the last boot parameters for this template, in XML format	read only
ha-always-run	True if an instance of this template will always restarted on another host in case of the failure of the host it is resident on	read/write
ha-restart-priority	1, 2, 3 or best effort. 1 is the highest restart priority	read/write
blobs	binary data store	read only
live	only relevant to a running VM.	read only

template-export

template-export template-uuid=<uuid_of_existing_template> filename=<filename_for_new_template>

Exports a copy of a specified template to a file with the specified new filename.

Update commands

Commands for working with updates to the OEM edition of XenServer. For commands relating to updating the standard non-OEM editions of XenServer, see [the section called “Patch \(update\) commands”](#) for details.

update-upload

update-upload file-name=<name_of_upload_file>

Streams a new software image to a OEM edition XenServer host. You must then restart the host for this to take effect.

User commands

user-password-change

user-password-change old=<old_password> new=<new_password>

Changes the password of the logged-in user. The old password field is not checked because you require supervisor privilege to make this call.

VBD commands

Commands for working with VBDs (Virtual Block Devices).

A VBD is a software object that connects a VM to the VDI, which represents the contents of the virtual disk. The VBD has the attributes which tie the VDI to the VM (is it bootable, its read/write metrics, and so on), while the VDI has the information on the physical attributes of the virtual disk (which type of SR, whether the disk is shareable, whether the media is read/write or read only, and so on).

The VBD objects can be listed with the standard object listing command (**xe vbd-list**), and the parameters manipulated with the standard parameter commands. See [the section called “Low-level param commands”](#) for details.

VBD parameters

VBDs have the following parameters:

Parameter Name	Description	Type
uuid	the unique identifier/object reference for the VBD	read only
vm-uuid	the unique identifier/object reference for the VM this VBD is attached to	read only
vm-name-label	the name of the VM this VBD is attached to	read only
vdi-uuid	the unique identifier/object reference for the VDI this VBD is mapped to	read only
vdi-name-label	the name of the VDI this VBD is mapped to	read only
empty	if true, this represents an empty drive	read only
device	the device seen by the guest, for example <code>hda1</code>	read only
userdevice	user-friendly device name	read/write
bootable	true if this VBD is bootable	read/write
mode	the mode the VBD should be mounted with	read/write
type	how the VBD appears to the VM, for example <code>disk</code> or <code>CD</code>	read/write
currently-attached	True if the VBD is currently attached on this host, false otherwise	read only

Parameter Name	Description	Type
storage-lock	True if a storage-level lock was acquired	read only
status-code	error/success code associated with the last attach operation	read only
status-detail	error/success information associated with the last attach operation status	read only
qos_algorithm_type	the QoS algorithm to use	read/write
qos_algorithm_params	parameters for the chosen QoS algorithm	read/write map parameter
qos_supported_algorithms	supported QoS algorithms for this VBD	read only set parameter
io_read_kbs	average read rate in kB per second for this VBD	read only
io_write_kbs	average write rate in kB per second for this VBD	read only
allowed-operations	list of the operations allowed in this state. This list is advisory only and the server state may have changed by the time this field is read by a client.	read only set parameter
current-operations	links each of the running tasks using this object (by reference) to a current_operation enum which describes the nature of the task.	read only set parameter
unpluggable	true if this VBD will support hot-unplug	read/write
attachable	True if the device can be attached	read only
other-config	additional configuration	read/write map parameter

vbd-create

```
vbd-create vm-uuid=<uuid_of_the_vm> device=<device_value>
vdi-uuid=<uuid_of_the_vdi_the_vbd_will_connect_to> [bootable=true] [type=<Disk | CD>] [mode=<RW | RO>]
```



Create a new VBD on a VM.

Appropriate values for the *device* field are listed in the parameter *allowed-VBD-devices* on the specified VM. Before any VBDs exist there, the allowable values are integers from 0-15.

If the *type* is *Disk*, *vdι-uuid* is required. Mode can be *RO* or *RW* for a Disk.

If the *type* is *CD*, *vdι-uuid* is optional; if no VDI is specified, an empty VBD will be created for the CD. Mode must be *RO* for a CD.

vbd-destroy

```
vbd-destroy uuid=<uuid_of_vbd>
```

Destroy the specified VBD.

If the VBD has its *other-config:owner* parameter set to *true*, the associated VDI will also be destroyed.

vbd-eject

```
vbd-eject uuid=<uuid_of_vbd>
```

Remove the media from the drive represented by a VBD. This command only works if the media is of a removable type (a physical CD or an ISO); otherwise an error message *VBD_NOT_REMOVABLE_MEDIA* is returned.

vbd-insert

```
vbd-insert uuid=<uuid_of_vbd> vdi-uuid=<uuid_of_vdi_containing_media>
```

Insert new media into the drive represented by a VBD. This command only works if the media is of a removable type (a physical CD or an ISO); otherwise an error message *VBD_NOT_REMOVABLE_MEDIA* is returned.

vbd-plug

```
vbd-plug uuid=<uuid_of_vbd>
```

Attempt to attach the VBD while the VM is in the running state.

vbd-unplug

```
vbd-unplug uuid=<uuid_of_vbd>
```

Attempts to detach the VBD from the VM while it is in the running state.

VDI commands

Commands for working with VDIs (Virtual Disk Images).

A VDI is a software object that represents the contents of the virtual disk seen by a VM, as opposed to the VBD, which is a connector object that ties a VM to the VDI. The VDI has the information on the physical attributes of the virtual disk (which type of SR, whether the disk is shareable, whether the media is read/write or read only, and so on), while the VBD has the attributes which tie the VDI to the VM (is it bootable, its read/write metrics, and so on).

The VDI objects can be listed with the standard object listing command (**xe vdi-list**), and the parameters manipulated with the standard parameter commands. See [the section called “Low-level param commands”](#) for details.

VDI parameters

VDIs have the following parameters:

Parameter Name	Description	Type
uuid	the unique identifier/object reference for the VDI	read only
name-label	the name of the VDI	read/write
name-description	the description string of the VDI	read/write
allowed-operations	a list of the operations allowed in this state	read only set parameter
current-operations	a list of the operations that are currently in progress on this VDI	read only set parameter
sr-uuid	SR in which the VDI resides	read only
vbd-uuids	a list of VBDs that refer to this VDI	read only set parameter
crashdump-uuids	list of crash dumps that refer to this VDI	read only set parameter
virtual-size	size of disk as presented to the VM, in bytes. Note that, depending on the storage backend type, the size may not be respected exactly	read only
physical-Utilization	amount of physical space that the VDI is currently taking up on the SR, in bytes	read only
type	type of VDI, for example, System or User	read only
sharable	true if this VDI may be shared	read only
read-only	true if this VDI can only be mounted read-only	read only
storage-lock	true if this VDI is locked at the storage level	read only
parent	references the parent VDI, if this VDI is part of a chain	read only
missing	true if SR scan operation reported this VDI as not present	read only
other-config	additional configuration information for this VDI	read/write map parameter
sr-name-label	name of the containing storage repository	read only
location	location information	read only

Parameter Name	Description	Type
managed	true if the VDI is managed	read only
xenstore-data	data to be inserted into the xenstore tree (/local/domain/0/backend/vbd/<domid>/<device-id>/sm-data) after the VDI is attached. This is generally set by the SM backends on vdi_attach .	read only map parameter
sm-config	SM dependent data	read only map parameter
is-a-snapshot	true if this VDI is a VM storage snapshot	read only
snapshot_of	the UUID of the storage this VDI is a snapshot of	read only
snapshots	the UUID(s) of all snapshots of this VDI	read only
snapshot_time	the timestamp of the snapshot operation that created this VDI	read only

vdi-clone

```
vdi-clone uuid=<uuid_of_the_vdi> [driver-params:<key=value>]
```

Create a new, writable copy of the specified VDI that can be used directly. It is a variant of **vdi-copy** that is capable of exposing high-speed image clone facilities where they exist.

The optional *driver-params* map parameter can be used for passing extra vendor-specific configuration information to the back end storage driver that the VDI is based on. See the storage vendor driver documentation for details.

vdi-copy

```
vdi-copy uuid=<uuid_of_the_vdi> sr-uuid=<uuid_of_the_destination_sr>
```

Copy a VDI to a specified SR.

vdi-create

```
vdi-create sr-uuid=<uuid_of_the_sr_where_you_want_to_create_the_vdi>
name-label=<name_for_the_vdi>
type=<system | user | suspend | crashdump>
virtual-size=<size_of_virtual_disk>
sm-config-*=<storage_specific_configuration_data>
```

Create a VDI.

The *virtual-size* parameter can be specified in bytes or using the IEC standard suffixes KiB (2¹⁰ bytes), MiB (2²⁰ bytes), GiB (2³⁰ bytes), and TiB (2⁴⁰ bytes).

Note:

SR types that support sparse allocation of disks (such as Local VHD and NFS) do not enforce virtual allocation of disks. Users should therefore take great care when over-allocating virtual disk space on an SR. If an over-



allocated SR does become full, disk space must be made available either on the SR target substrate or by deleting unused VDIs in the SR.

Note:

Some SR types might round up the *virtual-size* value to make it divisible by a configured block size.

vdi-destroy

```
vdi-destroy uuid=<uuid_of_vdi>
```

Destroy the specified VDI.

Note:

In the case of Local VHD and NFS SR types, disk space is not immediately released on **vdi-destroy**, but periodically during a storage repository scan operation. Users that need to force deleted disk space to be made available should call **sr-scan** manually.

vdi-forget

```
vdi-forget uuid=<uuid_of_vdi>
```

Unconditionally removes a VDI record from the database without touching the storage backend. In normal operation, you should be using **vdi-destroy** instead.

vdi-import

```
vdi-import uuid=<uuid_of_vdi> filename=<filename_of_raw_vdi>
```

Import a raw VDI.

vdi-introduce

```
vdi-introduce uuid=<uuid_of_vdi>  
sr-uuid=<uuid_of_sr_to_import_into>  
name-label=<name_of_the_new_vdi>  
type=<system | user | suspend | crashdump>  
location=<device_location_(varies_by_storage_type)>  
[name-description=<description_of_vdi>]  
[sharable=<yes | no>]  
[read-only=<yes | no>]  
[other-config=<map_to_store_misc_user_specific_data>]  
[xenstore-data=<map_to_of_additional_xenstore_keys>]  
[sm-config<storage_specific_configuration_data>]
```

Create a VDI object representing an existing storage device, without actually modifying or creating any storage. This command is primarily used internally to automatically introduce hot-plugged storage devices.

vdi-resize

```
vdi-resize uuid=<vdi_uuid> disk-size=<new_size_for_disk>
```

Resize the VDI specified by UUID.

vdi-snapshot

```
vdi-snapshot uuid=<uuid_of_the_vdi> [driver-params=<params>]
```



Produces a read-write version of a VDI that can be used as a reference for backup and/or template creation purposes. You can perform a backup from a snapshot rather than installing and running backup software inside the VM. The VM can continue running while external backup software streams the contents of the snapshot to the backup media. Similarly, a snapshot can be used as a "gold image" on which to base a template. A template can be made using any VDIs.

The optional *driver-params* map parameter can be used for passing extra vendor-specific configuration information to the back end storage driver that the VDI is based on. See the storage vendor driver documentation for details.

A clone of a snapshot should always produce a writable VDI.

vdi-unlock

```
vdi-unlock uuid=<uuid_of_vdi_to_unlock> [force=true]
```

Attempts to unlock the specified VDIs. If *force=true* is passed to the command, it will force the unlocking operation.

VIF commands

Commands for working with VIFs (Virtual network interfaces).

The VIF objects can be listed with the standard object listing command (**xe vif-list**), and the parameters manipulated with the standard parameter commands. See [the section called "Low-level param commands"](#) for details.

VIF parameters

VIFs have the following parameters:

Parameter Name	Description	Type
uuid	the unique identifier/object reference for the VIF	read only
vm-uuid	the unique identifier/object reference for the VM that this VIF resides on	read only
vm-name-label	the name of the VM that this VIF resides on	read only
allowed-operations	a list of the operations allowed in this state	read only set parameter
current-operations	a list of the operations that are currently in progress on this VIF	read only set parameter
device	integer label of this VIF, indicating the order in which VIF backends were created	read only
MAC	MAC address of VIF, as exposed to the VM	read only

Parameter Name	Description	Type
MTU	Maximum Transmission Unit of the VIF in bytes. This parameter is read-only, but you can override the MTU setting with the <i>mtu</i> key using the other-config <i>map</i> parameter. For example, to reset the MTU on a virtual NIC to use jumbo frames: <pre>xe vif-param-set \ uuid=<vif_uuid> \ other-config:mtu=9000</pre>	read only
currently-attached	true if the device is currently attached	read only
qos_algorithm_type	QoS algorithm to use	read/write
qos_algorithm_params	parameters for the chosen QoS algorithm	read/write map parameter
qos_supported_algorithms	supported QoS algorithms for this VIF	read only set parameter
MAC-autogenerated	True if the MAC address of the VIF was automatically generated	read only
other-config	Additional configuration key:value pairs	read/write map parameter
other-config:ethtool-rx	set to <i>on</i> to enable receive checksum, <i>off</i> to disable	read write
other-config:ethtool-tx	set to <i>on</i> to enable transmit checksum, <i>off</i> to disable	read write
other-config:ethtool-sg	set to <i>on</i> to enable scatter gather, <i>off</i> to disable	read write
other-config:ethtool-tso	set to <i>on</i> to enable tcp segmentation offload, <i>off</i> to disable	read write
other-config:ethtool-ufo	set to <i>on</i> to enable udp fragment offload, <i>off</i> to disable	read write
other-config:ethtool-gso	set to <i>on</i> to enable generic segmentation offload, <i>off</i> to disable	read write

Parameter Name	Description	Type
other-config:promiscuous	true to a VIF to be promiscuous on the bridge, so that it sees all traffic over the bridge. Useful for running an Intrusion Detection System (IDS) or similar in a VM.	read write
network-uuid	the unique identifier/object reference of the virtual network to which this VIF is connected	read only
network-name-label	the descriptive name of the virtual network to which this VIF is connected	read only
io_read_kbs	average read rate in kB/s for this VIF	read only
io_write_kbs	average write rate in kB/s for this VIF	read only

vif-create

```
vif-create vm-uuid=<uuid_of_the_vm> device=<see below>
network-uuid=<uuid_of_the_network_the_vif_will_connect_to> [mac=<mac_address>]
```

Create a new VIF on a VM.

Appropriate values for the *device* field are listed in the parameter *allowed-VIF-devices* on the specified VM. Before any VIFs exist there, the allowable values are integers from 0-15.

The *mac* parameter is the standard MAC address in the form aa:bb:cc:dd:ee:ff. If you leave it unspecified, an appropriate random MAC address will be created. You can also explicitly set a random MAC address by specifying *mac=random*.

vif-destroy

```
vif-destroy uuid=<uuid_of_vif>
```

Destroy a VIF.

vif-plug

```
vif-plug uuid=<uuid_of_vif>
```

Attempt to attach the VIF while the VM is in the running state.

vif-unplug

```
vif-unplug uuid=<uuid_of_vif>
```



Attempts to detach the VIF from the VM while it is running.

VLAN commands

Commands for working with VLANs (virtual networks). To list and edit virtual interfaces, refer to the PIF commands, which have a VLAN parameter to signal that they have an associated virtual network (see [the section called “PIF commands”](#)). For example, to list VLANs you need to use **xe pif-list**.

vlan-create

```
vlan-create pif-uuid=<uuid_of_pif> vlan=<vlan_number> network-uuid=<uuid_of_network>
```

Create a new VLAN on a XenServer host.

pool-vlan-create

```
vlan-create pif-uuid=<uuid_of_pif> vlan=<vlan_number> network-uuid=<uuid_of_network>
```

Create a new VLAN on all hosts on a pool, by determining which interface (for example, `eth0`) the specified network is on (on each host) and creating and plugging a new PIF object one each host accordingly.

vlan-destroy

```
vlan-destroy uuid=<uuid_of_pif_mapped_to_vlan>
```

Destroy a VLAN. Requires the UUID of the PIF that represents the VLAN.

VM commands

Commands for controlling VMs and their attributes.

VM selectors

Several of the commands listed here have a common mechanism for selecting one or more VMs on which to perform the operation. The simplest way is by supplying the argument `vm=<name_or_uuid>`. An easy way to get the uuid of an actual VM is to, for example, execute **xe vm-list power-state=running**. (The full list of fields that can be matched can be obtained by the command **xe vm-list params=all**.) For example, specifying `power-state=halted` will select all VMs whose `power-state` parameter is equal to `halted`. Where multiple VMs are matching, the option `--multiple` must be specified to perform the operation. The full list of parameters that can be matched is described at the beginning of this section, and can be obtained by the command `xe vm-list params=all`.

The VM objects can be listed with the standard object listing command (**xe vm-list**), and the parameters manipulated with the standard parameter commands. See [the section called “Low-level param commands”](#) for details.

VM parameters

VMs have the following parameters:

Note:

All writable VM parameter values can be changed while the VM is running, but the new parameters are *not* applied dynamically and will not be applied until the VM is rebooted.

Parameter Name	Description	Type
uuid	the unique identifier/object reference for the VM	read only
name-label	the name of the VM	read/write
name-description	the description string of the VM	read/write
user-version	string for creators of VMs and templates to put version information	read/write
is-a-template	False unless this is a template; template VMs can never be started, they are used only for cloning other VMs Note that setting <code>is-a-template</code> using the CLI is not supported.	read/write
is-control-domain	True if this is a control domain (domain 0 or a driver domain)	read only
power-state	current power state	read only
memory-dynamic-max	dynamic maximum in bytes	read/write
memory-dynamic-min	dynamic minimum in bytes	read/write
memory-static-max	statically-set (absolute) maximum in bytes. If you want to change this value, the VM must be shut down.	read/write
memory-static-min	statically-set (absolute) minimum in bytes. If you want to change this value, the VM must be shut down.	read/write
suspend-VDI-uuid	the VDI that a suspend image is stored on	read only

Parameter Name	Description	Type
VCPUs-params	<p>configuration parameters for the selected VCPU policy.</p> <p>You can tune a VCPU's pinning with</p> <pre>xe vm-param-set \ uuid=<vm_uuid> \ VCPUs-params:mask=1,2,3</pre> <p>The selected VM will then run on physical CPUs 1, 2, and 3 only.</p> <p>You can also tune the VCPU priority (xen scheduling) with the <i>cap</i> and <i>weight</i> parameters; for example</p> <pre>xe vm-param-set \ uuid=<template_uuid> \ VCPUs-params:weight=512 xe vm-param-set \ uuid=<template UUID> \ VCPUs-params:cap=100</pre> <p>A VM with a weight of 512 will get twice as much CPU as a domain with a weight of 256 on a contended XenServer host. Legal weights range from 1 to 65535 and the default is 256.</p> <p>The cap optionally fixes the maximum amount of CPU a VM will be able to consume, even if the XenServer host has idle CPU cycles. The cap is expressed in percentage of one physical CPU: 100 is 1 physical CPU, 50 is half a CPU, 400 is 4 CPUs, etc. The default, 0, means there is no upper cap.</p>	read/write map parameter
VCPUs-max	maximum number of virtual CPUs.	read/write
VCPUs-at-startup	boot number of virtual CPUs	read/write

Parameter Name	Description	Type
actions-after-crash	action to take if the VM crashes. For PV guests, valid parameters are: preserve (for analysis only), coredump_and_restart (record a coredump and reboot VM), coredump_and_destroy (record a coredump and leave VM halted), restart (no coredump and restart VM), and destroy (no coredump and leave VM halted).	read/write
console-uuids	virtual console devices	read only set parameter
platform	platform-specific configuration	read/write map parameter
allowed-operations	list of the operations allowed in this state	read only set parameter
current-operations	a list of the operations that are currently in progress on the VM	read only set parameter
allowed-VBD-devices	list of VBD identifiers available for use, represented by integers of the range 0-15. This list is informational only, and other devices may be used (but may not work).	read only set parameter
allowed-VIF-devices	list of VIF identifiers available for use, represented by integers of the range 0-15. This list is informational only, and other devices may be used (but may not work).	read only set parameter
HVM-boot-policy	the boot policy for HVM guests. Either BIOS Order or an empty string.	read/write
HVM-boot-params	the <i>order</i> key controls the HVM guest boot order, represented as a string where each character is a boot method: <i>d</i> for the CD/DVD, <i>c</i> for the root disk, and <i>n</i> for network PXE boot. The default is <i>dc</i> .	read/write map parameter

Parameter Name	Description	Type
HVM-shadow-multiplier	Floating point value which controls the amount of shadow memory overhead to grant the VM. Defaults to 1.0 (the minimum value), and should only be changed by advanced users.	read/write
PV-kernel	path to the kernel	read/write
PV-ramdisk	path to the initrd	read/write
PV-args	string of kernel command line arguments	read/write
PV-legacy-args	string of arguments to make legacy VMs boot	read/write
PV-bootloader	name of or path to bootloader	read/write
PV-bootloader-args	string of miscellaneous arguments for the bootloader	read/write
last-boot-CPU-flags	describes the CPU flags on which the VM was last booted	read only
resident-on	the XenServer host on which a VM is currently resident	read only
affinity	a XenServer host which the VM has preference for running on; used by the xe vm-start command to decide where to run the VM	read/write
other-config	<p>A list of key/value pairs that specify additional configuration parameters for the VM</p> <p>For example, a VM will be started automatically after host boot if the other-config parameter includes the key/value pair <i>auto_poweron: true</i></p>	read/write map parameter

Parameter Name	Description	Type
start-time	timestamp of the date and time that the metrics for the VM were read, in the form <code>yyyymmddThh:mm:ss z</code> , where <code>z</code> is the single-letter military timezone indicator, for example, <code>Z</code> for UTC (GMT)	read only
install-time	timestamp of the date and time that the metrics for the VM were read, in the form <code>yyyymmddThh:mm:ss z</code> , where <code>z</code> is the single-letter military timezone indicator, for example, <code>Z</code> for UTC (GMT)	read only
memory-actual	the actual memory being used by a VM	read only
VCPUs-number	<p>the number of virtual CPUs assigned to the VM</p> <p>For a paravirtualized Linux VM, this number can differ from <code>VCPUs-max</code> and can be changed without rebooting the VM using the vm-vcpu-hotplug command. See the section called “vm-vcpu-hotplug”. Windows VMs always run with the number of vCPUs set to <code>VCPUs-max</code> and must be rebooted to change this value.</p> <p>Note that performance will drop sharply if you set <code>VCPUs-number</code> to a value greater than the number of physical CPUs on the XenServer host.</p>	read only
VCPUs-Utilization	a list of virtual CPUs and their weight	read only map parameter
os-version	the version of the operating system for the VM	read only map parameter
PV-drivers-version	the versions of the paravirtualized drivers for the VM	read only map parameter

Parameter Name	Description	Type
PV-drivers-up-to-date	flag for latest version of the paravirtualized drivers for the VM	read only
memory	memory metrics reported by the agent on the VM	read only map parameter
disks	disk metrics reported by the agent on the VM	read only map parameter
networks	network metrics reported by the agent on the VM	read only map parameter
other	other metrics reported by the agent on the VM	read only map parameter
guest-metrics-last-updated	timestamp when the last write to these fields was performed by the in-guest agent, in the form <code>yyyymmddThh:mm:ssz</code> , where <code>z</code> is the single-letter military timezone indicator, for example, <code>Z</code> for UTC (GMT)	read only
actions-after-shutdown	action to take after the VM has shutdown	read/write
actions-after-reboot	action to take after the VM has rebooted	read/write
possible-hosts	potential hosts of this VM	read only
dom-id	domain ID (if available, -1 otherwise)	read only
recommendations	XML specification of recommended values and ranges for properties of this VM	read only
xenstore-data	data to be inserted into the xenstore tree (<code>/local/domain/<domid>/vm-data</code>) after the VM is created	read/write map parameter
is-a-snapshot	True if this VM is a snapshot	read only
snapshot_of	the UUID of the VM this is a snapshot of	read only
snapshots	the UUID(s) of all snapshots of this VM	read only

Parameter Name	Description	Type
snapshot_time	the timestamp of the snapshot operation that created this VM snapshot	read only
memory-target	the target amount of memory set for this VM	read only
blocked-operations	lists the operations that cannot be performed on this VM	read/write map parameter
last-boot-record	record of the last boot parameters for this template, in XML format	read only
ha-always-run	True if this VM will always restarted on another host in case of the failure of the host it is resident on	read/write
ha-restart-priority	1, 2, 3 or best effort. 1 is the highest restart priority	read/write
blobs	binary data store	read only
live	True if the VM is running, false if HA suspects that the VM may not be running.	read only

vm-cd-add

```
vm-cd-add cd-name=<name_of_new_cd> device=<integer_value_of_an_available_vbd>
[<vm-selector>=<vm_selector_value>...]
```

Add a new virtual CD to the selected VM. The *device* parameter should be selected from the value of the *allowed-VBD-devices* parameter of the VM.

The VM or VMs on which this operation should be performed are selected using the standard selection mechanism (see [VM selectors](#)). Optional arguments can be any number of the [VM parameters](#) listed at the beginning of this section.

vm-cd-eject

```
vm-cd-eject [<vm-selector>=<vm_selector_value>...]
```

Eject a CD from the virtual CD drive. This command only works if there is one and only one CD attached to the VM. When there are two or more CDs, use the command **xe vbd-eject** and specify the UUID of the VBD.

The VM or VMs on which this operation should be performed are selected using the standard selection mechanism (see [VM selectors](#)). Optional arguments can be any number of the [VM parameters](#) listed at the beginning of this section.

vm-cd-insert

```
vm-cd-insert cd-name=<name_of_cd> [<vm-selector>=<vm_selector_value>...]
```

Insert a CD into the virtual CD drive. This command will only work if there is one and only one empty CD device attached to the VM. When there are two or more empty CD devices, use the **xe vbd-insert** command and specify the UUIDs of the VBD and of the VDI to insert.

The VM or VMs on which this operation should be performed are selected using the standard selection mechanism (see [VM selectors](#)). Optional arguments can be any number of the [VM parameters](#) listed at the beginning of this section.

vm-cd-list

```
vm-cd-list [vbd-params] [vdi-params] [<vm-selector>=<vm_selector_value>...]
```

Lists CDs attached to the specified VMs.

The VM or VMs on which this operation should be performed are selected using the standard selection mechanism (see [VM selectors](#)). Optional arguments can be any number of the [VM parameters](#) listed at the beginning of this section.

You can also select which VBD and VDI parameters to list.

vm-cd-remove

```
vm-cd-remove cd-name=<name_of_cd> [<vm-selector>=<vm_selector_value>...]
```

Remove a virtual CD from the specified VMs.

The VM or VMs on which this operation should be performed are selected using the standard selection mechanism (see [VM selectors](#)). Optional arguments can be any number of the [VM parameters](#) listed at the beginning of this section.

vm-clone

```
vm-clone new-name-label=<name_for_clone>  
[new-name-description=<description_for_clone>] [<vm-selector>=<vm_selector_value>...]
```

Clone an existing VM, using storage-level fast disk clone operation where available. Specify the name and the optional description for the resulting cloned VM using the *new-name-label* and *new-name-description* arguments.

The VM or VMs on which this operation should be performed are selected using the standard selection mechanism (see [VM selectors](#)). Optional arguments can be any number of the [VM parameters](#) listed at the beginning of this section.

vm-compute-maximum-memory

```
vm-compute-maximum-memory total=<amount_of_available_physical_ram_in_bytes>  
[approximate=<add_overhead_memory_for_additional_vCPUS? true | false>  
[<vm-selector>=<vm_selector_value>...]
```

Calculate the maximum amount of static memory which can be allocated to an existing VM, using the total amount of physical RAM as an upper bound. The optional parameter *approximate* reserves sufficient extra memory in the calculation to account for adding extra vCPUs into the VM at a later date.



For example:

```
xe vm-compute-maximum-memory vm=testvm total=`xe host-list params=memory-free --minimal`
```

This command uses the value of the *memory-free* parameter returned by the **xe host-list** command to set the maximum memory of the VM named *testvm*.

The VM or VMs on which this operation will be performed are selected using the standard selection mechanism (see [VM selectors](#)). Optional arguments can be any number of the [VM parameters](#) listed at the beginning of this section.

vm-copy

```
vm-copy new-name-label=<name_for_copy> [new-name-description=<description_for_copy>]  
[sr-uuid=<uuid_of_sr>] [<vm-selector>=<vm_selector_value>...]
```

Copy an existing VM, but without using storage-level fast disk clone operation (even if this is available). The disk images of the copied VM are guaranteed to be "full images" - that is, not part of a copy-on-write (CoW) chain.

Specify the name and the optional description for the resulting copied VM using the *new-name-label* and *new-name-description* arguments.

Specify the destination SR for the resulting copied VM using the *sr-uuid*. If this parameter is not specified, the destination is the same SR that the original VM is in.

The VM or VMs on which this operation should be performed are selected using the standard selection mechanism (see [VM selectors](#)). Optional arguments can be any number of the [VM parameters](#) listed at the beginning of this section.

vm-crashdump-list

```
vm-crashdump-list [<vm-selector>=<vm selector value>...]
```

List crashdumps associated with the specified VMs.

If the optional argument *params* is used, the value of *params* is a string containing a list of parameters of this object that you want to display. Alternatively, you can use the keyword *all* to show all parameters. If *params* is not used, the returned list shows a default subset of all available parameters.

The VM or VMs on which this operation should be performed are selected using the standard selection mechanism (see [VM selectors](#)). Optional arguments can be any number of the [VM parameters](#) listed at the beginning of this section.

vm-data-source-forget

```
vm-data-source-forget data-source=<name_description_of_data-source> [<vm-selector>=<vm  
selector value>...]
```

Stop recording the specified data source for a VM, and forget all of the recorded data.

The VM or VMs on which this operation should be performed are selected using the standard selection mechanism (see [VM selectors](#)). Optional arguments can be any number of the [VM parameters](#) listed at the beginning of this section.

vm-data-source-list

```
vm-data-source-list [<vm-selector>=<vm selector value>...]
```




List the data sources that can be recorded for a VM.

The VM or VMs on which this operation should be performed are selected using the standard selection mechanism (see [VM selectors](#)). Optional arguments can be any number of the [VM parameters](#) listed at the beginning of this section.

vm-data-source-query

```
vm-data-source-query data-source=<name_description_of_data-source> [<vm-selector>=<vm selector value>...]
```

Display the specified data source for a VM.

The VM or VMs on which this operation should be performed are selected using the standard selection mechanism (see [VM selectors](#)). Optional arguments can be any number of the [VM parameters](#) listed at the beginning of this section.

vm-data-source-record

```
vm-data-source-record data-source=<name_description_of_data-source> [<vm-selector>=<vm selector value>...]
```

Record the specified data source for a VM.

This will write the information from the data source to the VM's persistent performance metrics database. This database is distinct from the normal agent database for performance reasons.

Data sources have the true/false parameters *standard* and *enabled*, which can be seen in the output of the **vm-data-source-list** command. If *enabled=true*, the data source metrics are currently being recorded to the performance database; if *enabled=false* they are not. Data sources with *standard=true* have *enabled=true* and have their metrics recorded to the performance database by default. Data sources which have *standard=false* have *enabled=false* by default. The **vm-data-source-record** command sets *enabled=false*.

Once enabled, you can stop recording the metrics of the data source by running the [vm-data-source-forget command](#).

The VM or VMs on which this operation should be performed are selected using the standard selection mechanism (see [VM selectors](#)). Optional arguments can be any number of the [VM parameters](#) listed at the beginning of this section.

vm-destroy

```
vm-destroy uuid=<uuid_of_vm>
```

Destroy the specified VM. This leaves the storage associated with the VM intact. To delete storage as well, use **xe vm-uninstall**.

vm-disk-add

```
vm-disk-add disk-size=<size_of_disk_to_add> device=<uuid_of_device> [<vm-selector>=<vm_selector_value>...]
```

Add a new disk to the specified VMs. Select the *device* parameter from the value of the *allowed-VBD-devices* parameter of the VMs.



The *disk-size* parameter can be specified in bytes or using the IEC standard suffixes KiB (2^{10} bytes), MiB (2^{20} bytes), GiB (2^{30} bytes), and TiB (2^{40} bytes).

The VM or VMs on which this operation should be performed are selected using the standard selection mechanism (see [VM selectors](#)). Optional arguments can be any number of the [VM parameters](#) listed at the beginning of this section.

vm-disk-list

```
vm-disk-list [vbd-params] [vdi-params] [<vm-selector>=<vm_selector_value>...]
```

Lists disks attached to the specified VMs. The *vbd-params* and *vdi-params* parameters control the fields of the respective objects to output and should be given as a comma-separated list, or the special key *all* for the complete list.

The VM or VMs on which this operation should be performed are selected using the standard selection mechanism (see [VM selectors](#)). Optional arguments can be any number of the [VM parameters](#) listed at the beginning of this section.

vm-disk-remove

```
vm-disk-remove device=<integer_label_of_disk> [<vm-selector>=<vm_selector_value>...]
```

Remove a disk from the specified VMs and destroy it.

The VM or VMs on which this operation should be performed are selected using the standard selection mechanism (see [VM selectors](#)). Optional arguments can be any number of the [VM parameters](#) listed at the beginning of this section.

vm-export

```
vm-export filename=<export_filename>  
[metadata=<true | false>]  
[<vm-selector>=<vm_selector_value>...]
```

Export the specified VMs (including disk images) to a file on the local machine. Specify the filename to export the VM into using the *filename* parameter. By convention, the filename should have a *.xva* extension.

If the *metadata* parameter is *true*, then the disks are not exported, and only the VM metadata is written to the output file. This is intended to be used when the underlying storage is transferred through other mechanisms, and permits the VM information to be recreated (see [the section called “vm-import”](#)).

The VM or VMs on which this operation should be performed are selected using the standard selection mechanism (see [VM selectors](#)). Optional arguments can be any number of the [VM parameters](#) listed at the beginning of this section.

vm-import

```
vm-import filename=<export_filename>  
[metadata=<true | false>]  
[preserve=<true | false>]  
[sr-uuid=<destination_sr_uuid>]
```

Import a VM from a previously-exported file. If *preserve* is set to *true*, the MAC address of the original VM will be preserved. The *sr-uuid* determines the destination SR to import the VM into, and is the default SR if not specified.



The *filename* parameter can also point to an XVA-format VM, which is the legacy export format from XenServer 3.2 and is used by some third-party vendors to provide virtual appliances. This format uses a directory to store the VM data, so set *filename* to the root directory of the XVA export and not an actual file. Subsequent exports of the imported legacy guest will automatically be upgraded to the new filename-based format, which stores much more data about the configuration of the VM.

Note:

The older directory-based XVA format does not fully preserve all the VM attributes. In particular, imported VMs will not have any virtual network interfaces attached by default. If networking is required, create one using **vif-create** and **vif-plug**.

If the *metadata* is *true*, then a previously exported set of metadata can be imported without their associated disk blocks. Metadata-only import will fail if any VDIs cannot be found (named by SR and *VDI.location*) unless the *--force* option is specified, in which case the import will proceed regardless. If disks can be mirrored or moved out-of-band then metadata import/export represents a fast way of moving VMs between disjoint pools (e.g. as part of a disaster recovery plan).

Note:

Multiple VM imports will be performed faster in serial than in parallel.

vm-install

```
vm-install new-name-label=<name>
[ template-uuid=<uuid_of_desired_template> | [template=<uuid_or_name_of_desired_template>]]
[ sr-uuid=<sr_uuid> | sr-name-label=<name_of_sr> ]
[ copy-bios-strings-from=<uuid of host> ]
```

Install a VM from a template. Specify the template name using either the *template-uuid* or *template* argument. Specify an SR other than the default SR using either the *sr-uuid* or *sr-name-label* argument. Specify to install BIOS-locked media using the *copy-bios-strings-from* argument.

vm-memory-shadow-multiplier-set

```
vm-memory-shadow-multiplier-set [<vm-selector>=<vm_selector_value>...]
[multiplier=<float_memory_multiplier>]
```

Set the shadow memory multiplier for the specified VM.

This is an advanced option which modifies the amount of *shadow memory* assigned to a hardware-assisted VM. In some specialized application workloads, such as Citrix XenApp, extra shadow memory is required to achieve full performance.

This memory is considered to be an overhead. It is separated from the normal memory calculations for accounting memory to a VM. When this command is invoked, the amount of free XenServer host memory will decrease according to the multiplier, and the *HVM_shadow_multiplier* field will be updated with the actual value which Xen has assigned to the VM. If there is not enough XenServer host memory free, then an error will be returned.

The VMs on which this operation should be performed are selected using the standard selection mechanism (see [VM selectors](#) for more information).

vm-migrate

```
vm-migrate [[host-uuid=<destination XenServer host UUID> ] | [host=<name or UUID of destination XenServer host> ]] [<vm-selector>=<vm_selector_value>...] [live=<true | false>]
```



Migrate the specified VMs between physical hosts. The *host* parameter can be either the name or the UUID of the XenServer host.

By default, the VM will be suspended, migrated, and resumed on the other host. The *live* parameter activates XenMotion and keeps the VM running while performing the migration, thus minimizing VM downtime to less than a second. In some circumstances such as extremely memory-heavy workloads in the VM, XenMotion automatically falls back into the default mode and suspends the VM for a brief period of time before completing the memory transfer.

The VM or VMs on which this operation should be performed are selected using the standard selection mechanism (see [VM selectors](#)). Optional arguments can be any number of the [VM parameters](#) listed at the beginning of this section.

vm-reboot

```
vm-reboot [<vm-selector>=<vm_selector_value>...] [force=<true>]
```

Reboot the specified VMs.

The VM or VMs on which this operation should be performed are selected using the standard selection mechanism (see [VM selectors](#)). Optional arguments can be any number of the [VM parameters](#) listed at the beginning of this section.

Use the *force* argument to cause an ungraceful shutdown, akin to pulling the plug on a physical server.

vm-reset-powerstate

```
vm-reset-powerstate [<vm-selector>=<vm_selector_value>...] {force=true}
```

The VM or VMs on which this operation should be performed are selected using the standard selection mechanism (see [VM selectors](#)). Optional arguments can be any number of the [VM parameters](#) listed at the beginning of this section.

This is an *advanced* command only to be used when a member host in a pool goes down. You can use this command to force the pool master to reset the power-state of the VMs to be *halted*. Essentially this forces the lock on the VM and its disks so it can be subsequently started on another pool host. This call *requires* the force flag to be specified, and fails if it is not on the command-line.

vm-resume

```
vm-resume [<vm-selector>=<vm_selector_value>...] [force=<true | false>] [on=<XenServer host UUID>]
```

Resume the specified VMs.

The VM or VMs on which this operation should be performed are selected using the standard selection mechanism (see [VM selectors](#)). Optional arguments can be any number of the [VM parameters](#) listed at the beginning of this section.

If the VM is on a shared SR in a pool of hosts, use the *on* argument to specify which host in the pool on which to start it. By default the system will determine an appropriate host, which might be any of the members of the pool.

vm-shutdown

```
vm-shutdown [<vm-selector>=<vm_selector_value>...] [force=<true | false>]
```



Shut down the specified VM.

The VM or VMs on which this operation should be performed are selected using the standard selection mechanism (see [VM selectors](#)). Optional arguments can be any number of the [VM parameters](#) listed at the beginning of this section.

Use the *force* argument to cause an ungraceful shutdown, similar to pulling the plug on a physical server.

vm-start

```
vm-start [<vm-selector>=<vm_selector_value>...] [force=<true | false>] [on=<XenServer host UUID>] [--multiple]
```

Start the specified VMs.

The VM or VMs on which this operation should be performed are selected using the standard selection mechanism (see [VM selectors](#)). Optional arguments can be any number of the [VM parameters](#) listed at the beginning of this section.

If the VMs are on a shared SR in a pool of hosts, use the *on* argument to specify which host in the pool on which to start the VMs. By default the system will determine an appropriate host, which might be any of the members of the pool.

vm-suspend

```
vm-suspend [<vm-selector>=<vm_selector_value>...]
```

Suspend the specified VM.

The VM or VMs on which this operation should be performed are selected using the standard selection mechanism (see [VM selectors](#)). Optional arguments can be any number of the [VM parameters](#) listed at the beginning of this section.

vm-uninstall

```
vm-uninstall [<vm-selector>=<vm_selector_value>...] [force=<true | false>]
```

Uninstall a VM, destroying its disks (those VDI's that are marked RW and connected to this VM only) as well as its metadata record. To simply destroy the VM metadata, use **xe vm-destroy**.

The VM or VMs on which this operation should be performed are selected using the standard selection mechanism (see [VM selectors](#)). Optional arguments can be any number of the [VM parameters](#) listed at the beginning of this section.

vm-vcpu-hotplug

```
vm-vcpu-hotplug new-vcpus=<new_vcpu_count> [<vm-selector>=<vm_selector_value>...]
```

Dynamically adjust the number of VCPUs available to a running paravirtual Linux VM within the number bounded by the parameter *VCPUS-max*. Windows VMs always run with the number of VCPUs set to *VCPUS-max* and must be rebooted to change this value.

The paravirtualized Linux VM or VMs on which this operation should be performed are selected using the standard selection mechanism (see [VM selectors](#)). Optional arguments can be any number of the [VM parameters](#) listed at the beginning of this section.



vm-vif-list

```
vm-vif-list [<vm-selector>=<vm_selector_value>...]
```

Lists the VIFs from the specified VMs.

The VM or VMs on which this operation should be performed are selected using the standard selection mechanism (see [VM selectors](#)). Note that the selectors operate on the VM records when filtering, and *not* on the VIF values. Optional arguments can be any number of the [VM parameters](#) listed at the beginning of this section.

Workload Balancing commands

Commands for controlling the Workload Balancing feature.

pool-initialize-wlb

```
pool-initialize-wlb wlb_url=<wlb_server_address> \  
wlb_username=<wlb_server_username> \  
wlb_password=<wlb_server_password> \  
xenserver_username=<pool_master_username> \  
xenserver_password=<pool_master_password>
```

Starts the Workload Balancing service on a pool.

Note:

Initializing a pool requires running two commands. After executing the **pool-initialize-wlb** command, execute **xe pool-param-set wlb-enabled=true uuid=<pool-uuid>[]**.

pool-param-set other-config

Use the **pool-param-set other-config** command to specify the timeout when communicating with the WLB server. All requests are serialized, and the timeout covers the time from a request being queued to its response being completed. In other words, slow calls cause subsequent ones to be slow. Defaults to 30 seconds if unspecified or if the setting cannot be parsed.

```
xe pool-param-set other-config:wlb_timeout=<0.01> \  
uuid=<315688af-5741-cc4d-9046-3b9cea716f69>
```

host-retrieve-wlb-evacuate-recommendations

```
host-retrieve-wlb-evacuate-recommendations uuid=<host_uuid>
```

Returns the evacuation recommendations for a host, and a reference to the UUID of the recommendations object.

vm-retrieve-wlb-recommendations

Returns the workload balancing recommendations for the selected VM. The simplest way to select the VM on which the operation is to be performed is by supplying the argument `vm=<name_or_uuid>`. VMs can also be specified by filtering the full list of VMs on the values of fields. For example, specifying `power-state=halted` selects all VMs whose power-state is halted. Where multiple VMs are matching, specify the option `--multiple` to perform the operation. The full list of fields that can be matched can be obtained by the command **xe vm-list params=all**. If no parameters to select VMs are given, the operation will be performed on all VMs.



pool-certificate-list

XenServer's Workload Balancing component lets you use certificates to secure communication between XenServer pools and the Workload Balancing server. You can either use the default certificate, which is created automatically during Workload Balancing installation, or you can specify a certificate you have already created.

To use your own certificate, the certificate must be in X.509 format. If you want to import the certificate into XenServer's certificate store, during Workload Balancing installation you must specify for WLB to use an existing certificate and export it. Currently, you must do this installing WLB with the Msiexec commands.

However, following installation you need to export the certificate again. When you export the certificate from Workload Balancing, Workload Balancing exports it in Base64 encoded format. You must convert the exported certificate into a Privacy Enhanced Mail (PEM) file or a .crt format by exporting it from Windows using Windows certificate management features so that XenServer can import it.

Note:

To convert the exported certificate into a PEM (.pem) file, copy it to your XenServer pool master and run the following commands:

```
openssl enc -base64 -in <exported_cert_name.crt> -out <certificate.pem>
```

After converting the certificate into .pem or .crt, you must load the certificate onto servers across the pool by doing the following:

1. List any existing certificates on the pool (using **xe pool-certificate-list**).
2. Install the certificate you specified during WLB installation (using **pool-certificate-install**).
3. Synchronize the certificate on all hosts in the pool (using **pool-certificate-sync**).
4. (Optional.) Instruct XenServer to require a certificate before connecting (using **pool-certificate-sync**).

```
pool-certificate-list
```

Lists all installed SSL certificates.

pool-certificate-install

```
pool-certificate-install filename=<certificatefilename>
```

Run this command on the pool to install the certificate you specified during WLB installation on the pool master. Before installing the certificate on the master, it must be exported in either .pem or .crt format. If you are exporting the certificate using Windows certificate management features, select the Base 64 encoded X.509 format.

Typically, when you installed WLB, you may have named the certificate something like wlbcert.cer. Simply renaming the file wlbcert.crt is not sufficient. You must export the certificate so the file formats is converted into a format XenServer is expecting to receive.

pool-certificate-sync

```
pool-certificate-install
```

Run this command on the pool, after running the pool-certificate-install command, to make sure the certificate and certificate revocation lists are synchronized from the pool master to all slaves on the pool.



pool-param-set

pool-param-set wlb-verify-cert=<true> uuid=<uuid_of_pool>

Run this command on the pool, after running the pool-certificate-sync command, to make XenServer always verify the certificate when communicating with the Workload Balancing server.

Tip:

Pressing the **Tab** key automatically populates the UUID of the pool.

pool-deconfigure-wlb

Permanently deletes all workload balancing configuration.

pool-retrieve-wlb-configuration

Prints all workload balancing configuration to standard out.

pool-retrieve-wlb-recommendations

Prints all workload balancing recommendations to standard out.

pool-retrieve-wlb-report

Gets a WLB report of the specified type and saves it to the specified file. The available reports are:

- pool_health
- pool_audit_history poolid
- pool_optimization_history
- host_health_history
- optimization_performance_history
- pool_health_history
- vm_movement_history
- vm_performance_history

Example usage for each report type is shown below. The *utcoffset* parameter specifies the number of hours ahead or behind of UTC for your time zone. The *start* parameter and *end* parameters specify the number of hours to report about. For example specifying *start=-3* and *end=0* will cause WLB to report on the last 3 hour of activity.

```
xe pool-retrieve-wlb-report report=pool_health \  
poolid=<51e411f1-62f4-e462-f1ed-97c626703cae> \  
utcoffset=<-5> \  
start=<-3> \  
end=<0> \  
filename=</pool_health.txt>
```

```
xe pool-retrieve-wlb-report report=host_health_history \  
hostid=<e26685cd-1789-4f90-8e47-a4fd0509b4a4> \  
utcoffset=<-5> \  
start=<-3> \  
end=<0> \  
filename=</host_health_history.txt>
```




```
xe pool-retrieve-wlb-report report=optimization_performance_history \  
poolid=<51e411f1-62f4-e462-f1ed-97c626703cae> \  
utcoffset=<-5> \  
start=<-3> \  
end=<0> \  
filename=</optimization_performance_history.txt>  
  
xe pool-retrieve-wlb-report report=pool_health_history \  
poolid=<51e411f1-62f4-e462-f1ed-97c626703cae> \  
utcoffset=<-5> \  
start=<-3> \  
end=<0> \  
<filename=/pool_health_history.txt>  
  
xe pool-retrieve-wlb-report report=vm_movement_history \  
poolid=<51e411f1-62f4-e462-f1ed-97c626703cae> \  
utcoffset=<-5> \  
start=<-5> \  
end=<0> \  
filename=</vm_movement_history.txt>  
  
xe pool-retrieve-wlb-report report=vm_performance_history \  
hostid=<e26685cd-1789-4f90-8e47-a4fd0509b4a4> \  
utcoffset=<-5> \  
start=<-3> \  
end=<0> \  
<filename=/vm_performance_history.txt>
```

pool-send-wlb-configuration

Modifies Workload Balancing configuration settings, including thresholds, WLB power-management settings, and weightings. It is not mandatory to configure all settings with the command. You can configure only some parameters, but not all, if desired.

Before using the pool-send-wlb-configuration command, learn about the default values on your system by running pool-retrieve-wlb-configuration.

When you run the pool-retrieve-wlb-configuration command, additional parameters appear that are not documented in this section. Citrix does not recommend editing these parameters.

```
pool-send-wlb-configuration  
[  
    config:HostMemoryThresholdCritical=<HostCpuThresholdCritical=value>  
    config:HostMemoryThresholdHigh=<HostMemoryThresholdHigh=value>config:HostPifReadThresholdCritical=<HostPifReadThresholdCritical=value>  
    config:HostPifReadThresholdHigh=<HostPifReadThresholdHigh=value>  
    config:set_host_configuration=<true | false> ...]
```

Use the **pool-send-wlb-configuration** command with the arguments **<ParticipatesInPowerManagement>** and **<set_host_configuration>** to configure Workload Balancing's Host Power Management feature.

```
xe pool-send-wlb-configuration \  
config:<host_21_> \  
ParticipatesInPowerManagement=<true> \  
config:set_host_configuration=<true>
```

Index

A

AMD-V (AMD hardware virtualization), 19

C

CD commands, *xe* CLI, 148

CLI (see command line interface)

Command line interface (CLI)

- basic *xe* syntax, 142

- Bonding commands, 147

- CD commands, 148

- command types, 144

- console commands, 149

- event commands, 150

- host (XenServer host) commands, 151

- log commands, 161

- low-level list commands, 146

- low-level parameter commands, 145

- message commands, 162

- network commands, 163

- overview,

- parameter types, 145

- patch commands, 164

- PBD commands, 165

- PIF commands, 166

- Resource pool commands, 170

- shorthand *xe* syntax, 143

- special characters and syntax, 143

- Storage Manager commands, 173

- Storage repository (SR) commands, 174

- task commands, 176

- Template commands, 177

- update commands, 185

- user commands, 185

- VBD commands, 186

- VDI commands, 188

- VIF commands, 192

- VLAN commands, 195

- VM commands, 195

- xe* command reference,

Console commands, *xe* CLI, 149

Constraints on XenServer hosts joining resource pool, 19

Creating a resource pool, 20

E

Event commands, *xe* CLI, 150

F

Fibre Channel storage area network (SAN), 53

Filer, NetApp, 43

FlexVol, NetApp, 43

H

Hardware virtualization

- AMD-V, 19

- Intel VT, 19

HBA (see Host bus adapter)

Host (XenServer host) commands, *xe* CLI, 151

Host bus adapter, 53

I

Intel VT (Intel hardware virtualization), 19

iSCSI, 47

L

Log commands, *xe* CLI, 161

Logical Volume Management (LVM), 40, 41

Logs, XenServer host, 140

M

Machine failures in a resource pool, 133

Message commands, *xe* CLI, 162

N

NAS (see Network attached storage (NFS))

NetApp Filer, 43

Network attached storage (NFS), 52

Network bonding commands, *xe* CLI, 147

Network commands, *xe* CLI, 163

Networking VMs,

Networking XenServer hosts

- Initial configuration after installation, 72

P

Patch commands, *xe* CLI, 164

PBD commands, *xe* CLI, 165

PIF commands, *xe* CLI, 166

Pool commands, *xe* CLI, 170

Q

QoS settings

- virtual disk, 61

R

Removing XenServer host from a resource pool, 22

Requirements, for creating resource pools, 19

Resource pool,

- constraints on XenServer hosts joining, 19

- coping with machine failures, 133

- creating, 20

- master, 19, 133, 133

- member, 133, 133
- removing XenServer host from, 22
- requirements for creating, 19

S

- SAN (see Storage Area Network)
- Shared network attached storage (NFS), 52
- Shared storage, 21
- Storage Area Network, 47
- Storage Manager commands, *xe* CLI, 173
- Storage repository (SR)
 - CD-ROM, 41
 - commands, *xe* CLI, 174
 - DVD-ROM, 41
 - Fibre Channel storage area network (SAN), 53
 - local disk, 40
 - local hotplug devices, 41
 - NetApp Filer, 43
 - overview, 30
 - shared iSCSI storage area network (SAN), 47
 - shared network attached storage (NFS), 52
 - USB read/write device, 41

T

- Task commands, *xe* CLI, 176
- Template commands, *xe* CLI, 177
- Troubleshooting
 - xen-bugtool*, ,
 - XenCenter logs, 141
 - XenServer host logs, 140, 140

U

- Update commands, *xe* CLI, 185
- User commands, *xe* CLI, 185

V

- VBD commands, *xe* CLI, 186
- VDI commands, *xe* CLI, 188
- VIF (virtual interface), 70
- VIF commands, *xe* CLI, 192
- Virtual network, 70
- VLAN commands, *xe* CLI, 195
- VM
 - commands, *xe* CLI, 195
 - networking,
 - Virtual disk QoS settings, 61
- VT (Intel hardware virtualization), 19

X

- xe* command reference,
 - basic *xe* syntax, 142
 - Bonding commands, 147
 - CD commands, 148

- command types, 144
- console commands, 149
- event commands, 150
- host (XenServer host) commands, 151
- log commands, 161
- low-level list commands, 146
- low-level parameter commands, 145
- message commands, 162
- network commands, 163
- parameter types, 145
- patch commands, 164
- PBD commands, 165
- PIF commands, 166
- Resource pool commands, 170
- shorthand *xe* syntax, 143
- special characters and syntax, 143
- Storage Manager commands, 173
- Storage repository (SR) commands, 174
- task commands, 176
- Template commands, 177
- update commands, 185
- user commands, 185
- VBD commands, 186
- VDI commands, 188
- VIF commands, 192
- VLAN commands, 195
- VM commands, 195
- xe* commands, command line interface (CLI),
 - xen-bugtool*, , 140
- XenServer host
 - constraints, on joining resource pool, 19
 - joining in resource pools,
 - logs, 140
 - networks,
 - requirements, for joining resource pools, 19
- XenServer host troubleshooting
 - logs, 140
 - xen-bugtool*, , 140